



Snapshots and Clones with Oracle on Linux

Implementation Guide

Overview of the steps and considerations for creating and deploying snapshots and clones with Oracle Database running on Linux platforms.

Contents

- 1. Introduction..... 3**
 - 1.1. Terminology 3
 - 1.2. FSP: Flash Storage Platform..... 3
- 2. Preparing the Source..... 4**
 - 2.1. Considerations 4
 - 2.1.1. Volume mapping..... 4
 - 2.1.2. LUN based snapshots..... 4
 - 2.2. Install SDM and Configure Storage Connectivity..... 4
 - 2.2.1. Install the SDM and Oracle agent 5
 - 2.2.2. Connect the SDM to the storage tier..... 5
 - 2.2.3. Verify the SDM configuration in Symphony 6
 - 2.2.4. Set the overriddenrive parameter 7
 - 2.3. Configure the Oracle Agent..... 7
 - 2.3.1. Test Database Connectivity..... 7
 - 2.3.2. Configure the desired Database 8
 - 2.3.3. Add the Database storage mounts to the agent config..... 9
- 3. Creating and Scheduling Snapshots 9**
 - 3.1. Snapshot Resource Area (SRA) Considerations 9
 - 3.2. Snapshot Group 10
 - 3.2.1. Add Snapshot Group..... 10
 - 3.2.2. Create Snapshot Policy for the Snapshot Group..... 10
 - 3.2.2.1. Create Snapshot Resource form 11
 - 3.2.2.2. Snapshot Resource Policy..... 12
 - 3.2.2.3. Create Snapshot Policy form 13
 - 3.2.2.4. Assign the Snapshot Group to the SAN Client 15
 - 3.3. Test Snapshots 16
- 4. Create Clones..... 19**
 - 4.1. Choose Snapshot Source 19
 - 4.2. Choose Clone Type..... 19
 - 4.3. Assign Clones to Host..... 21
- 5. Mounting Clones to a Linux Host 22**
 - 5.1. Prepare Volumes..... 22
- 6. Attaching Cloned Databases to Oracle..... 23**
- 7. Remap a Cloned Database..... 26**
 - 7.1. Prepare the Database and Volumes 26
 - 7.2. Remap the Clone..... 27
 - 7.3. Online the Diskgroup and Database 29



1. Introduction

Snapshots and clones can have many uses ranging from performance scaling to workload isolation to deployment simplicity to data security. This makes snapshots and clones one of the most powerful tools on modern enterprise storage platforms. The following is a guide that walks the reader through snapshot and clone creation and deployment as well as covers some of the common considerations used for decision making. The guide should be viewed as an overview of the steps, in order, that should be taken. This guide does not replace the best practices guide or the array's user guides which should be used as companion documents.

1.1. Terminology

Before we get into the steps, let's clarify terminology.

- Snapshots. Point in time markers saved on the storage array. These markers allow an administrator to revert a LUN, or group of LUNs, back to the state of the data at the time of the marker.
- Clone. A new LUN, or set of LUNs, created to contain data with a state as of a point in time (a snapshot).
- Thin clones are merely pointers to the existing LUN or LUNs whereby changes to the data are held in the newly created clone LUNs but all of the original, non-changed data is still stored in the source LUN.
- Thick clones are fully hydrated LUNs whereby the array will make an exact copy of the LUN into new LUNs, and these new LUNs will act independently. Once a thick clone is created it cannot be remapped to another snapshot (point-in-time marker).
- Remap. The ability to change the source point-in-time marker (snapshot) for a thin clone.
- Rollback. Ability to revert the state of the LUN or LUNs back to a previous point in time marker (snapshot).

1.2. FSP: Flash Storage Platform

Violin Memory's Flash Storage Platform (FSP) is the combination of custom-for-flash hardware and custom-for-flash software delivering both cost effective and performance based storage solutions. Concerto OS7 is the software layer of the FSP that delivers a full suite of data services including protection, efficiency, continuity, migration and scaling. These features can also be combined for increased value and security. For example, an application consistent snapshot can reside on mirrored LUNs or the source LUNs can also be replicated to remote sites, etc.

Also included in the FSP are host based agents and tools which allow for advanced application consistent snapshots as well as other features. On the Linux platform, these tools and agents include:

- SDM. The SAN Disk Manager (SDM) is the foundation configuration tool which allows an administrator to connect the host to the FSP, configure Fibre channel or iSCSI connectivity, get information about FSP devices presented, manage the other host agents and much more.
- Snapshot Agent for Oracle Database. The snapshot process is initiated and owned by the FSP. This agent receives communication from the FSP during a snapshot and coordinates with the Oracle Database to place the database into backup



mode for the duration of the snapshot, delivering an application consistent snapshot. This agent also allows for configuration of any scripts that may need to be ran before and/or after the snapshot.

2. Preparing the Source

The snapshot and clone process journey starts with the source host. In this section we will review the considerations and pre-snapshot steps that should be taken before creating and scheduling snapshots.

2.1. Considerations

Snapshots can be used for many purposes so considerations can cover a large amount of ground. The following are a few things to keep in mind when designing the architecture of the primary host.

2.1.1. Volume mapping

As each LUN and clone of a LUN receives a unique identifier (UUID), there is no concern over presenting either a LUN or a clone of that LUN to another or the same Linux host. However, as the contents of a LUN and its clone will be *exactly* the same initially, care should be taken if the discovery of a second, identical storage device could confuse an application. For example, Oracle ASM stamps each device under its control with a device header. Detecting a duplicate header on a second device on the same system could lead to undefined behavior by Oracle ASM. Plan carefully how LUNs and their clones will be presented to and consumed by hosts.

2.1.2. LUN based snapshots

Concerto OS7 is a LUN-based snapshot technology. This means that while the array can, for application consistent snapshots, coordinate with the Oracle Database to put the database into backup mode during the brief snapshot process for the entire LUN or LUN group. Clones will therefore be of the entire LUN or LUN group, not specific volumes or files on the LUN. So, if the intent of the architecture is to have the ability to only clone (or restore) one database, then multiple databases should not be resident on the same LUN as the cloned LUN will contain all volumes and files (databases). Similarly, the administrator must decide whether to include Oracle binaries on the same LUN(s) as the Oracle Database files, causing these to be snapped together. In some installations, it is preferable to place these on separate LUNs, allowing for infrequent or manual snaps of the binaries that do not change frequently and may not be desired to clone or replicate along with the database files. For deployments utilizing Oracle ASM, all LUNs in a diskgroup must be included in a snapshot for that snapshot to have value. Snapshot groups are used for such grouping, making administration easier and protecting write ordering within the group.

2.2. Install SDM and Configure Storage Connectivity

The SAN Disk Manager (SDM) is the management tool that allows for a variety of tasks to be accomplished on the host server. With it an administrator can create LUNs, set up connections to storage, review agent events, review FSP storage device attributes. This is a required piece of software to be installed on any host that will be the source of snapshots, as this tool is configures the connection to the FSP, configures the Fibre Channel or iSCSI connectivity settings, and manages the host agents after install. The SDM is command line only on Linux.

```

root@Oracle-1:~
[root@Oracle-1 ~]# ll ~/IMA-Linux-2.51-514-Violin.run
-rwxr-xr-x. 1 root root 5146035 Dec 11 13:09 /root/IMA-Linux-2.51-514-Violin.run
[root@Oracle-1 ~]# ll ~/sanora-4.50-700.i386.rpm
-rw-r--r--. 1 root root 1928395 Aug 19 14:20 /root/sanora-4.50-700.i386.rpm
[root@Oracle-1 ~]#

```

2.2.1. Install the SDM and Oracle agent

Refer to the FSP User Guide for full install instructions. The following is an example only.

```

[root@Oracle-1 /]# ~/IMA-Linux-2.51-514-Violin.run
Creating directory IMA-Linux-2.51-514-Violin
Verifying archive integrity... All good.
Uncompressing IMA.....
Installing IMA for Linux, please wait...
Preparing... ##### [100%]
 1:al ##### [100%]
Preparing... ##### [100%]
 1:IMA ##### [100%]
Initialize IMA configuration file ..... [ OK ]
Starting IMA service ..... [ OK ]

[root@Oracle-1 /]# rpm --nodeps -ivh ~/sanora-4.50-700.i386.rpm
Preparing... ##### [100%]
 1:sanora ##### [100%]

Snapshot Agent for Oracle installation completed.

Please add the following line in /etc/profile or /etc/csh.login:
    export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:<OraLibPath> #ISORAENV
where <OraLibPath> is the path to the Oracle shared libraries.

Please log out, log back in as root, and run the following script to configure the agent:
    /usr/local/sanclient/agents/oracle/bin/ora_setup.sh
[root@Oracle-1 /]# vi /root/.bash_profile
[root@Oracle-1 /]# su -
[root@Oracle-1 ~]# set |grep LD_
LD_LIBRARY_PATH=/u01/app/oracle/product/11.2.0.4/rdbms/lib

```

2.2.2. Connect the SDM to the storage tier

If no FSPs are configured, then add one with the iscmcli addserv command:

```
root@Oracle-1:~  
[root@Oracle-1 ~]# iscmcli addserv -s 192.168.0.53 -u vmemadmin -c password -p 2 -w 11-22-33-44-55-66-77-88  
Begin to add server ...  
Server added successfully.  
[root@Oracle-1 ~]#
```

If an FSP has previously been configured, review its configuration via the `iscmcli listserv` command:

```
root@Oracle-1:~  
[root@Oracle-1 ~]# iscmcli listserv  
Total number of storage servers: 1  
Storage server: Prod-Array-2-mg-a  
IP : 192.168.0.53  
Server type : SAN  
Version : 7.1362 (Build 8983)  
Login User : vmemadmin  
Protocol : Fiber Channel  
<FC Settings>  
WWPN1 : 11-22-33-44-55-66-77-88  
  
[root@Oracle-1 ~]#
```

2.2.3. Verify the SDM configuration in Symphony

Within Symphony where the FSP is managed, click on the Manage tab, then the SAN Clients subtab. Locate the Linux host just registered and confirm it contains the login name specified in the previous step as part of the client name field. This is required for proper functioning of the agents later. If the client was previously added from within Symphony and does not show the login name in this field, you must delete the client registration and perform it as described in the previous step. Note that a SAN client may not be deleted while any LUN(s) are assigned to it, so unassign the LUN(s) before attempting to delete the client.

LUNS	SAN CLIENTS	SNAPSHOT RESOURCES	SNAPSHOT POLICIES	SNAPSHOTS	REPLICATIONS	SNAPSHOT GROUPS	STORAGE POOLS	PHYSICAL DEVICES	HBA's	
	ACCESS-5[vmemadmin]	Prod-Array-2-mg-a	192.168.0.30	Winnt4				<input checked="" type="checkbox"/>	<input type="checkbox"/>	00:11:22:33:44:55:66:77
	LD-Appliance	Prod-Array-2-mg-a	192.168.0.8	Linux	LD-Test-Dedup LD-Test-Thick	0 1	Read/Write (Shared) Read/Write (Shared)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	21:00:00:0E:1E:09:98:DF 21:00:00:0E:1E:09:98:DF
	ORACLE-1[vmemadmin]	Prod-Array-2-mg-a	192.168.0.101	Linux	Oracle-1TB-Dedupe-2 Oracle-1TB-Thick-1 Oracle-1TB-Thick-2 Oracle-1TB-Thick-3 Oracle-1TB-Thick-4	18 0 1 2 3	Read/Write (Shared) Read/Write (Shared) Read/Write (Shared) Read/Write (Shared) Read/Write (Shared)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	11:22:33:44:55:66:77:88

2.2.4. Set the overriddenrive parameter

The overriddenrive setting tells the IMA a file system it should look to include in a snapshot. While this is not specifically required for the Oracle agent to function properly, the setting must be set to a directory that exists on the system. Below is an example that will accomplish this on a standard Linux host:

```
[root@Oracle-1 ~]# ismcli getoverridedrive
There is no override drive setting.
```

Set the override using the following command, then review the new setting:

```
[root@Oracle-1 ~]# ismcli setoverridedrive -d "/usr"
Set override drive list successfully.
[root@Oracle-1 ~]# ismcli getoverridedrive
The current overridedrive is /usr
```

2.3. Configure the Oracle Agent

The Oracle agent is the process that coordinates with the Oracle Database to place the database into backup mode during a snapshot. This is all that is required for an application consistent snapshot of Oracle, as, while in backup mode, the database records both the before and after image of every changed block in the redo logs. This constitutes enough information for the database to correct any fractured blocks recorded in the snapshot. An administrator can review the logs of the Oracle Database instance (known as the instance's *alert log*) and see the database going into and out of backup mode.

The Oracle agent must be configured for each database separately, as it tracks the Oracle Home path and login credentials for each database. The database must be reachable via a TNS alias established in the RDBMS Oracle Home. Test connectivity to the database prior to attempting to configure the agent.

2.3.1. Test Database Connectivity

```
[root@Oracle-1 ~]# tnsping perf

TNS Ping Utility for Linux: Version 11.2.0.4.0 - Production on 16-AUG-2015 13:00:00

Copyright (c) 1997, 2013, Oracle. All rights reserved.

Used parameter files:
```

```
Used TNSNAMES adapter to resolve the alias
Attempting to contact (DESCRIPTION = (ADDRESS_LIST = (ADDRESS = (PROTOCOL = TCP)(Host =
Oracle-1)(Port = 1521))) (CONNECT_DATA = (SID = perf)))
OK (0 msec)
```

2.3.2. Configure the desired Database

Use the agent setup script to add the database to the agent configuration.

```
[root@Oracle-1 bin]# /usr/local/sanclient/agents/oracle/bin/ora_setup.sh

Oracle Agent Setup

Enter 1 for Oracle 11 agent setup
Enter 2 for Oracle 10 agent setup

Enter choice: 1
Snapshot agent for Oracle 11 setup
Do you want to setup the 64-bit Oracle 11 agent (y/n)? y
What is the Oracle lib path for shared libraries? (libclntsh.so.11.1)
Enter path: /app/oracle/product/11.2.0/dbhome_1/lib

Please make sure the Oracle lib path you entered above
is added to the LD_LIBRARY_PATH and LIBPATH environment variables

The following is the current LD_LIBRARY_PATH:
/app/oracle/product/11.2.0/dbhome_1/lib

If the current values displayed above do not contain the Oracle lib path,
please make sure the environment variables are configured correctly before
proceeding with the setup for Oracle agent.

If necessary, you may need to log out and log back in for newly added
environment variables to take effect.

Would you like to continue the setup (y/n)? y
Snapshot Agent Setup for Oracle 11 64 bit
If there are multiple Oracle instances on this system, please make sure all the
instances you plan to add to the agent configuration are running.

Is Oracle started? (y/n) y
What is the Oracle SID? perf
Please make sure Oracle instance perf is running.

What is the Oracle home path for SID perf? /app/oracle/product/11.2.0/dbhome_1
What is the Oracle admin user id for SID orcl? (system)
What is the Oracle admin password for SID orcl? oracle
What is the login name of the UNIX user account that owns Oracle software? oracle

Would you like to enter another Oracle SID? (y/n) n
Updating /etc/profile .....

-----
Environment variables updated. Please log out, log back in as root,
and restart the IMA / SAN Client service.
-----
```




2.3.3. Add the Database storage mounts to the agent config

Edit in the file system mount point(s) and ASM diskgroups where the database datafiles, controlfiles, and online redo logs are stored.

```
[root@Oracle-1 bin]# vi /usr/local/sanclient/agents/oracle11/etc/agt.ini
```

```

root@Oracle-1:/usr/local/fsima/log
[mount point list begin]
#enter a list of snapshot mount points (e.g. /mnt/disk1)
#/mnt/disk1
#/mnt/disk2
#.....
/usr
+DATA
+PERF
[mount point list end]

#enter a pre snapshot script file (e.g. /tmp/presnap.sh)
#presnap_script=/tmp/presnap.sh

#enter a post snapshot script file (e.g. /tmp/postsnap.sh)
#postsnap_script=/tmp/postsnap.sh

```

3. Creating and Scheduling Snapshots

Moving from the database host to Symphony, now it is time to create the actual snapshot resources, policies and schedules. It is important to note that once initial configurations are set on the host, all management and monitoring should be done on the FSP via Symphony. Concerto OS7, the array software, is the overall owner of the snapshot execution and initiates coordination with the host side agents. Symphony is the management console tool that allows for configuration and monitoring of the Concerto OS7 software.

3.1. Snapshot Resource Area (SRA) Considerations

Snapshots track the application changes by saving both the new and old blocks. The older versions of the blocks are saved into a Snapshot Resource Area (SRA) that allows an administrator to allocate specific amounts of space for this function based upon their system's requirements for data security and availability. The SRA can auto grow and shrink and has many different retention policy options.


If any or all of the LUNs in a snapshot group will have unique retention or sizing requirements then the SRA per LUN should be set up individually. This is done via the SNAPSHOT RESOURCES tab in Symphony. If, however, the LUNs can have the same sizing and retention policies then a batch process can be utilized for the SRA creation via the SNAPSHOT GROUP tab in Symphony. It is most common, and easier, to use the SNAPSHOT GROUPS tab's wizard to create the SRA in a batch for all the LUNs at the same time. It is a best practice to utilize a few common policies over the snapshot use cases instead of unique policies for each, unless

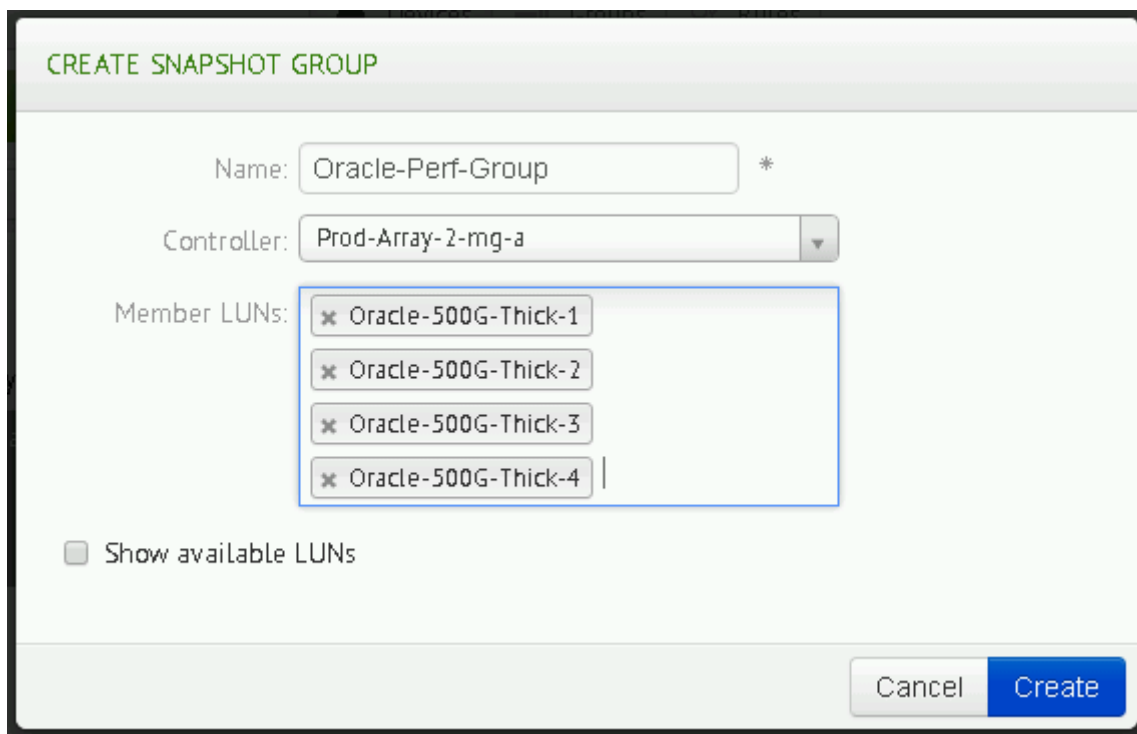
separate policies are required by the particular implementation.

3.2. Snapshot Group


With the wizard in the SNAPSHOT GROUPS tab, all of the snapshot resources, policies and scheduling can be created.

3.2.1. Add Snapshot Group

The first step is to create a new Snapshot Group. In the SNAPSHOTS GROUPS tab, click on the plus sign icon  and choose an appropriate name. This is also where you will choose all of the LUNs to include in the group and the controller to coordinate the snapshot process.



3.2.2. Create Snapshot Policy for the Snapshot Group

Next, start the wizard by clicking on the camera icon  and choosing Create Snapshot Policy. This wizard will take you through the steps required to set up the resource area (SRA), retention policy and scheduling. Click Start on the first page of the wizard to move to the next step.



CREATE SNAPSHOT POLICY

Snapshot Resource needs to be created before **Snapshot Policy** can be created.

LUN: Oracle-500G-Thick-2, Oracle-500G-Thick-4, Oracle-500G-Thick-1, Oracle-500G-Thick-3

Step 1: Add Snapshot Resource

Step 2: Add Snapshot Policy

Click the **Start** button to create **Snapshot Resource** followed by **Snapshot Policy**.

Cancel Start

3.2.2.1. Create Snapshot Resource form

In this form you will configure several items. Fill out the Storage Policy, Resource Policy and Advanced tabs. Do not fill out the Reclamation Policy tab yet. Click on the Resource Policy tab after defining the storage policy on the first tab.

CREATE SNAPSHOT RESOURCE

Storage Policy
Resource Policy
Advanced
Reclamation Policy

LUN: Oracle-500G-Thick-2, Oracle-500G-Thick-4, Oracle-500G-Thick-1, Oracle-500G-Thick-3

Storage Pool: A Storage-Pool *

Allocated Size: 10 GiB *

Available Size: 19130 GiB (Storage Pool Size 40960 GiB)

Threshold: 75 % *

Expand Automatically: ON

Increment Size: 20 % *

Max Size: 200 GiB

Cancel
Create

3.2.2.2. Snapshot Resource Policy

On this tab, you define how snapshots are kept and their priority. The options are as follows:

Preserve all snapshots: grow the SRA as needed and potentially stall writes to the LUN(s) if the SRA is full and cannot grow

Preserve recent snapshots: selecting this will cause a new dialog to pop asking how many snapshots to preserve. These will be kept in rolling fashion, with the SRA growing as needed to accommodate this number of snaps.

Maintain writes: grow the SRA as needed until it cannot grow anymore, then purge snapshots, starting with the oldest, to free space to keep from stalling incoming writes to LUN(s)

The screenshot shows a web-based configuration form titled "CREATE SNAPSHOT RESOURCE". The form has four tabs: "Storage Policy", "Resource Policy", "Advanced", and "Reclamation Policy". The "Resource Policy" tab is currently selected. Under this tab, there are three radio button options for "Resource Policy": "Preserve All Snapshots", "Preserve Recent Snapshots", and "Always Maintain Write Operations". The "Always Maintain Write Operations" option is selected. Below these options is a toggle switch for "Enable Snapshot Notification", which is currently set to "ON". At the bottom right of the form, there are two buttons: "Cancel" and "Create".

3.2.2.3. Create Snapshot Policy form

In this form, fill out the Schedule, Notification and Retention tabs. Skip the other two tabs.

When setting up the schedule, make sure to leave time before the first snapshot is taken to finish up any other configurations that may need to be done.

CREATE SNAPSHOT POLICY

Schedule
Notification
Replication
Retention
CDP

Group: Oracle-Perf-Group

Schedule Snapshots: ON

Initial Snapshot On: *

Snapshot Frequency: *

Trigger Snapshot Notification: ON

Snapshot Notification Every: scheduled snapshots *

Maximum Snapshots: *

Note that the application consistent or “notification” process requires communicating with the host-side agent to snap each configured database. If there is more than one database being snapshotted, this could add to the length of time the whole snapshot process takes. Note also that the setting specified on this tab is stored as a default for this snapshot policy. Manual snapshots may override this, and snapshots collected via the Schedule mechanism may choose to employ application notification for every snapshot or once every N snapshots, depending on business needs.

As shown here, administrators may be very specific in the snapshot retention policy.

CREATE SNAPSHOT POLICY

Schedule
Notification
Replication
Retention
CDP

Keep: All Snapshots
 Most Recent Snapshots
 Snapshots based on following Rules

All for past: *


Hourly from the past: days, closest to min *

Daily from the past: days, closest to hour *

Weekly from the past: weeks, closest to *

Monthly from the past: months, closest to day *

3.2.2.4. Assign the Snapshot Group to the SAN Client

Locate the Computer icon  while still on the SNAPSHOT GROUPS tab and with the newly created Snapshot Group selected, and click it and choose Assign to SAN Client from the drop-down menu.

LUNs	SAN CLIENTS	SNAPSHOT RESOURCES	SNAPSHOT POLICIES	SNAPSHOTS	REPLICATIONS	SNAPSHOT GROUPS	STORAGE POOLS
<div style="border: 1px solid gray; padding: 5px;"> + ✎ 🗑 ↺ 📄 ↻ </div>							
<div style="border: 1px solid gray; padding: 5px;"> ➔ Assign to SAN Client ➔ Unassign from SAN Client </div>							
Name	Type	Controller	Members	Clients	Snapshot Enabled	CDP Enabled	CDP Journal Storage Pool Snapshot Notification Enabled
Oracle-Perf-Group	Snapshot	Prod-Array-2-mg-a	Oracle-500G-Thick-2 Oracle-500G-Thick-4 Oracle-500G-Thick-1 Oracle-500G-Thick-3	Prod-ESXi-1	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
Oracle1-ASM-Data	Snapshot	Prod-Array-2-mg-a	Oracle-1TB-Thick-1 Oracle-1TB-Thick-3	ORACLE-1[vmemadmin] Prod-ESXi-1	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Choose the appropriate Linux host, and click Assign. Note that this is the same as assigning each of the LUNs in the group to the SAN client.

ASSIGN LUNS

Select LUNs to Assign

Select All Deselect All

Search

LUN	LUN #	Type	Group
<input checked="" type="checkbox"/> Oracle-500G-Thick-1	auto	Thick	Oracle-Perf-Group
<input checked="" type="checkbox"/> Oracle-500G-Thick-2	auto	Thick	Oracle-Perf-Group
<input checked="" type="checkbox"/> Oracle-500G-Thick-3	auto	Thick	Oracle-Perf-Group
<input checked="" type="checkbox"/> Oracle-500G-Thick-4	auto	Thick	Oracle-Perf-Group

Showing 1 to 4 of 4 entries (filtered from 31 total entries)

Available LUN #: 4-17, 19-255

Access:

ALL Targets: ON

Assign to SAN Clients

Select All Deselect All

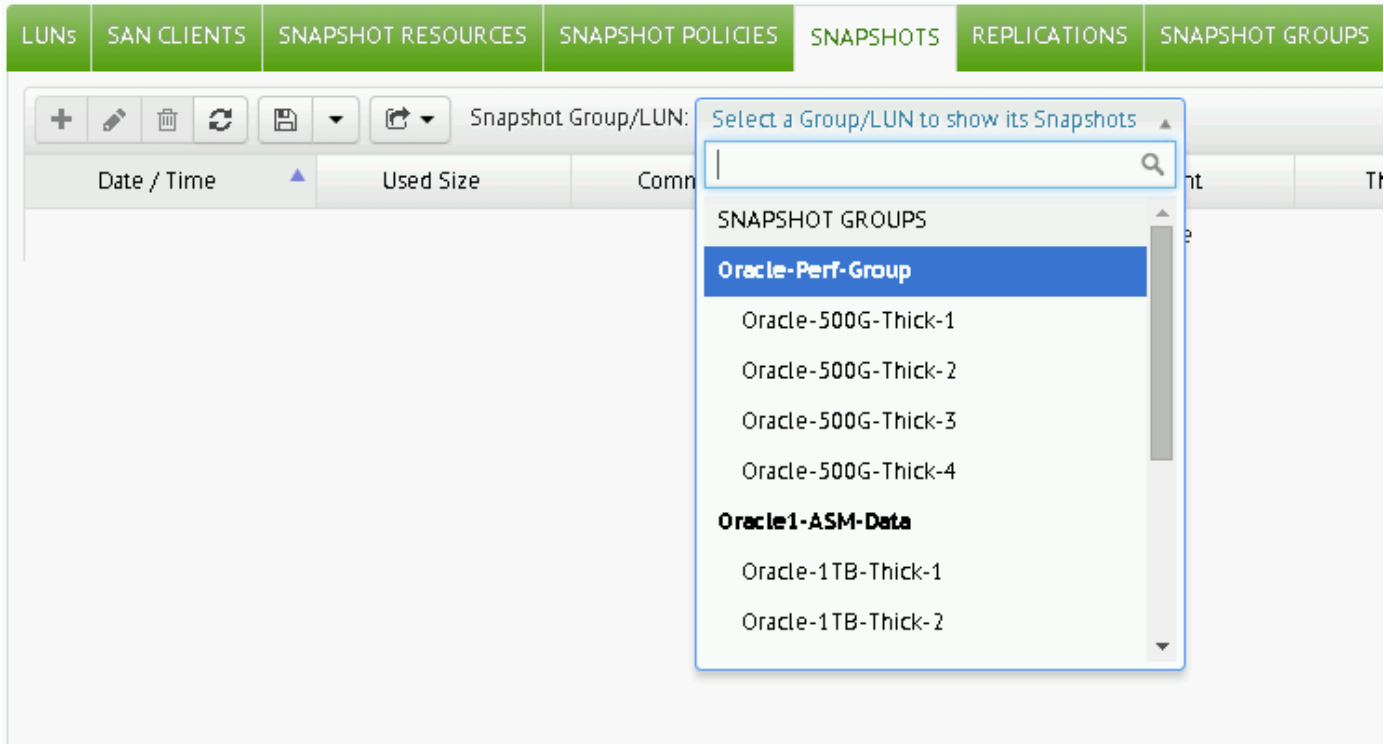
Search

Client	Allowed
<input type="checkbox"/> ACCESS-5[vmemadmin]	<input checked="" type="checkbox"/>
<input type="checkbox"/> LD-Appliance	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> ORACLE-1[vmemadmin]	<input checked="" type="checkbox"/>
<input type="checkbox"/> Prod-ESXi-1	<input checked="" type="checkbox"/>
<input type="checkbox"/> WindowsOracleServ	<input checked="" type="checkbox"/>

Showing 1 to 5 of 5 entries

3.3. Test Snapshots

The process of creating and scheduling snapshots is now done. You can review the snapshots as they are created by the system by going to the SNAPSHOTS tab in Symphony. Use the “Snapshot Group/LUN” drop down to review the existing snapshots for individual LUNs or snapshot groups.



Ensure the database is in ARCHIVELOG mode, as backup mode is meaningless in NOARCHIVELOG mode:

```

oracle@Oracle-1:~
SQL> archive log list;
Database log mode          No Archive Mode
Automatic archival        Disabled
Archive destination        USE_DB_RECOVERY_FILE_DEST
Oldest online log sequence 1282
Current log sequence       1284
SQL> shutdown immediate;
Database closed.
Database dismounted.
ORACLE instance shut down.
SQL> startup mount
ORACLE instance started.


Total System Global Area 3841671168 bytes
Fixed Size                2258960 bytes
Variable Size             2248148976 bytes
Database Buffers         1577058304 bytes
Redo Buffers              14204928 bytes
Database mounted.
SQL> alter database archivelog;

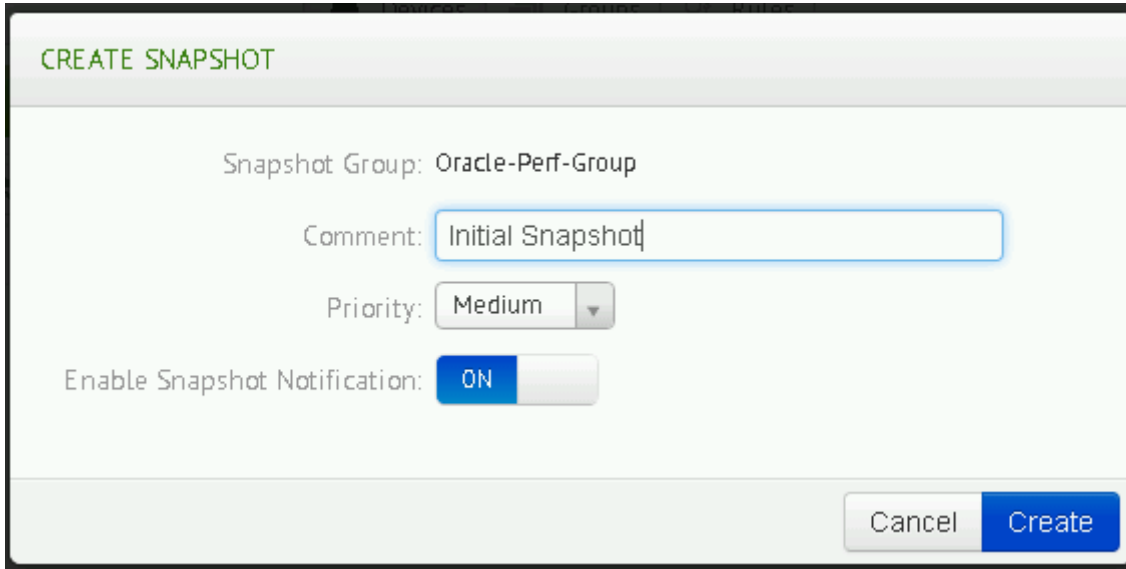
Database altered.

SQL> alter database open;

Database altered.

SQL>
    
```

Click the plus sign  to create a new snapshot manually, and specify that application notification be included.



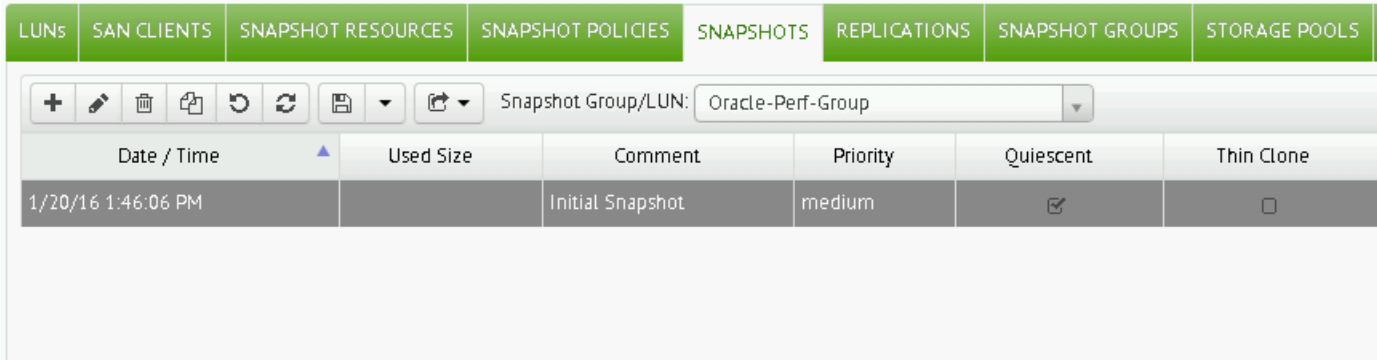
On the Linux host, locate the database instance alert log for the database previously configured with the Oracle agent configuration utility, and view the log contents immediately after the snapshot process completes and Symphony lists your new snapshot in its inventory.

```

oracle@Oracle-1:~
Wed Jan 20 13:46:11 2016
ALTER database backup controlfile to '/app/oracle/diag/rdbms/perf/perf/trace/is_oract100.bak'
Completed: ALTER database backup controlfile to '/app/oracle/diag/rdbms/perf/perf/trace/is_oract100.bak'
alter database backup controlfile to trace
Backup controlfile written to trace file /app/oracle/diag/rdbms/perf/perf/trace/perf_ora_12064.trc
Completed: alter database backup controlfile to trace
ALTER tablespace "SYSTEM" begin backup
Completed: ALTER tablespace "SYSTEM" begin backup
ALTER tablespace "SYSAUX" begin backup
Completed: ALTER tablespace "SYSAUX" begin backup
ALTER tablespace "UNDOTBS1" begin backup
Completed: ALTER tablespace "UNDOTBS1" begin backup
ALTER tablespace "USERS" begin backup
Completed: ALTER tablespace "USERS" begin backup
ALTER tablespace "TPCCTAB" begin backup
Completed: ALTER tablespace "TPCCTAB" begin backup
ALTER tablespace "SLOB" begin backup
Completed: ALTER tablespace "SLOB" begin backup
Wed Jan 20 13:46:18 2016
ALTER tablespace "SYSTEM" end backup
Completed: ALTER tablespace "SYSTEM" end backup
ALTER tablespace "SYSAUX" end backup
Completed: ALTER tablespace "SYSAUX" end backup
ALTER tablespace "UNDOTBS1" end backup
Completed: ALTER tablespace "UNDOTBS1" end backup
ALTER tablespace "USERS" end backup
Completed: ALTER tablespace "USERS" end backup
ALTER tablespace "TPCCTAB" end backup
Completed: ALTER tablespace "TPCCTAB" end backup
ALTER tablespace "SLOB" end backup
Completed: ALTER tablespace "SLOB" end backup
ALTER SYSTEM ARCHIVE LOG
Wed Jan 20 13:46:18 2016
Thread 1 advanced to log sequence 1288 (LGWR switch)
  Current log# 1 seq# 1288 mem# 0: +PERF/perf/online/seq/group_1.258.899799879
Archived Log entry 4 added for thread 1 sequence 1287 ID 0x491f51c3 dest 1:

```


In the resulting grid, the Quiescent column indicates which snapshots are application consistent (checked) or crash consistent (not checked).

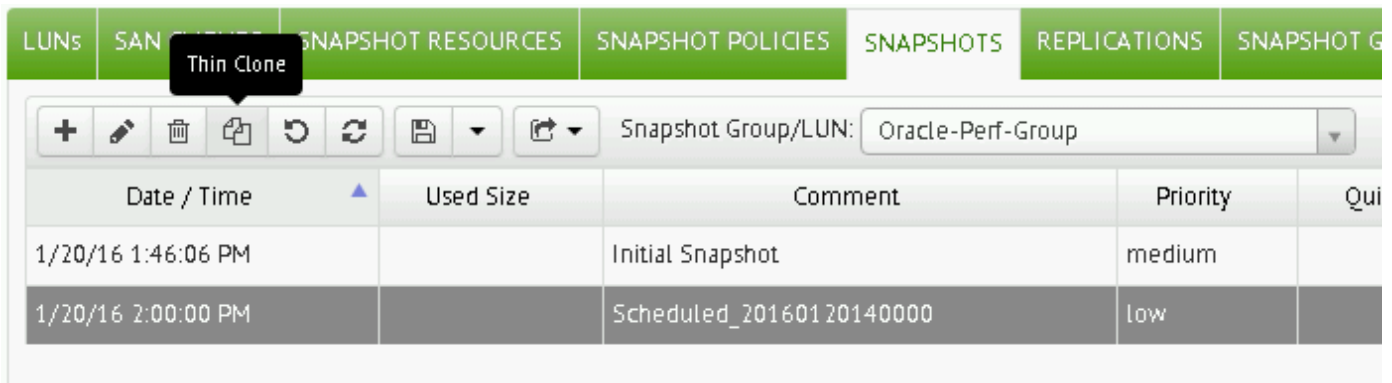


4. Create Clones

Now that snapshots are being taken, you can choose to turn one LUN or snapshot group into a thick or thin clone and mount it back to the same host or a different host.

4.1. Choose Snapshot Source

Choose one of the LUNs or snapshot groups in the “Snapshot Group/LUN” drop down. Choose a row item that matches your requirements. Then click the “Thin Clone” button  from the task bar.



4.2. Choose Clone Type

The default is a thin clone, which does not require any additional space initially. To create a fully hydrated thick clone as of the time of the selected snapshot, use the Instant Copy button on the bottom of the Thin Clone Snapshot form. This will start a background process that will copy data into the destination LUNs, resulting in a full separate copy of the source LUNs while allowing you to start using the LUNs immediately.

THIN CLONE SNAPSHOT

Snapshot | Storage Policy

Group: Oracle-Perf-Group

Snapshot At: 1/20/16 2:00:00 PM - Scheduled_20160120140000 *

Storage Pool: A Storage-Pool *

Allocated Size: 5 GiB *

Thin Clone Name Format: {n}_{t} ? *

Thin Clone Names: Oracle-500G-Thick-2_20160120140000 Oracle-500G-Thick-4_20160120140000
Oracle-500G-Thick-1_20160120140000 Oracle-500G-Thick-3_20160120140000

Instant Copy: OFF

Cancel Clone


Within seconds, the thin clones show as online.

LUNs																	
SAN CLIENTS		SNAPSHOT RESOURCES		SNAPSHOT POLICIES		SNAPSHOTS		REPLICATIONS		SNAPSHOT GROUPS		STORAGE POOLS		PHYSICAL DEVICES		HBAs	
Oracle-500G																	
LUN	LUN Type	Controller	Storage Pool	Snapshot Group	Tags	Status	Total Size	Allocated Size	Used Size	Used %	Clients						
Oracle-500G-Thick-1	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-1_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									
Oracle-500G-Thick-2	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-2_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									
Oracle-500G-Thick-3	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-3_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									
Oracle-500G-Thick-4	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-4_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									

4.3. Assign Clones to Host

Choose which host to assign the thin clone. This can be the same host as the source database or an entirely new host server.

Start by navigating to the LUNs tab in Symphony. Then highlight one of the new LUNs created by the thin clone process and click

the computer button  in the task bar. Then choose "Assign to SAN Client"

LUNs																	
SAN CLIENTS		SNAPSHOT RESOURCES		SNAPSHOT POLICIES		SNAPSHOTS		REPLICATIONS		SNAPSHOT GROUPS		STORAGE POOLS		PHYSICAL DEVICES		HBAs	
Oracle-500G																	
LUN	LUN Type	Controller	Storage Pool	Snapshot Group	Tags	Status	Total Size	Allocated Size	Used Size	Used %	Clients						
Oracle-500G-Thick-1	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-1_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									
Oracle-500G-Thick-2	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-2_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									
Oracle-500G-Thick-3	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-3_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									
Oracle-500G-Thick-4	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB			ORACLE-1[vmemadmin] Prod-ESXI-1						
Oracle-500G-Thick-4_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB									

In the ASSIGN LUNS form, select all of the LUNs in the snapshot group on the left and the desired host on the right and click the Assign button at the bottom.

ASSIGN LUNS

Select LUNs to Assign

Select All Deselect All

Oracle-500G-Thick

LUN	LUN #	Type	Grp
<input checked="" type="checkbox"/> Oracle-500G-Thick-2_20160120140000	auto	Thin Clone	
<input type="checkbox"/> Oracle-500G-Thick-3		Thick	Oracle
<input checked="" type="checkbox"/> Oracle-500G-Thick-3_20160120140000	auto	Thin Clone	
<input type="checkbox"/> Oracle-500G-Thick-4		Thick	Oracle
<input checked="" type="checkbox"/> Oracle-500G-Thick-4_20160120140000	auto	Thin Clone	

Showing 4 to 8 of 8 entries (filtered from 31 total entries)

Assign to SAN Clients

Select All Deselect All

Search

Client	Allowed
<input type="checkbox"/> ACCESS-5[vmemadmin]	<input checked="" type="checkbox"/>
<input type="checkbox"/> LD-Appliance	<input checked="" type="checkbox"/>
<input type="checkbox"/> ORACLE-1[vmemadmin]	<input checked="" type="checkbox"/>
<input checked="" type="checkbox"/> Oracle-2	<input checked="" type="checkbox"/>
<input type="checkbox"/> Prod-ESXI-1	<input checked="" type="checkbox"/>

Showing 1 to 5 of 5 entries

Available LUN #: 0-255

Access: Read/Write (Shared)

All Targets: ON

Cancel Assign

5. Mounting Clones to a Linux Host

Moving to the destination host for the clone, this section covers how to configure the newly assigned LUN or LUNs.

5.1. Prepare Volumes

If the new clone is being assigned to a new host for the first time, then Linux will have no mapping of the LUN. If using a physical Linux host, rescan the SCSI bus via `rescan-scsi-bus.sh` or by some other means. Once the LUN(s) are recognized by Linux, multipath aliases may be created based on device serial numbers, if desired. This is not necessary in the case of ASM volumes, as ASM will search available devices for the ASM header stamp and online the diskgroup if all members are present.

Example command sequence:

```
[root@Oracle-2 ~]# /usr/bin/rescan-scsi-bus.sh
...
<output truncated>
...
0 new or changed device(s) found.
```

```

0 remapped or resized device(s) found.
0 device(s) removed.
[root@oracle-2 ~]# multipath -F
[root@oracle-2 ~]# multipath -v2
...
<output truncated>
...
create: SviolIN_CONCERTO_ARRAY_7XYIHVSV7UVY undef VIOLIN,CONCERTO ARRAY
size=500G features='1 queue_if_no_path' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=130 status=undef
  |- 2:0:3:4 sde 8:64 undef ready running
  |- 2:0:1:4 sdo 8:224 undef ready running
  |- 1:0:0:4 sdy 65:128 undef ready running
  `-- 1:0:3:4 sdai 66:32 undef ready running
create: SviolIN_CONCERTO_ARRAY_ROQZ182CKK9H undef VIOLIN,CONCERTO ARRAY
size=500G features='1 queue_if_no_path' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=130 status=undef
  |- 2:0:3:5 sdf 8:80 undef ready running
  |- 2:0:1:5 sdp 8:240 undef ready running
  |- 1:0:0:5 sdz 65:144 undef ready running
  `-- 1:0:3:5 sdaj 66:48 undef ready running
create: SviolIN_CONCERTO_ARRAY_SAUSBK9MAP70 undef VIOLIN,CONCERTO ARRAY
size=500G features='1 queue_if_no_path' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=130 status=undef
  |- 2:0:3:6 sdg 8:96 undef ready running
  |- 2:0:1:6 sdq 65:0 undef ready running
  |- 1:0:0:6 sdaa 65:160 undef ready running
  `-- 1:0:3:6 sdak 66:64 undef ready running
create: SviolIN_CONCERTO_ARRAY_S667E400MH6R undef VIOLIN,CONCERTO ARRAY
size=500G features='1 queue_if_no_path' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=130 status=undef
  |- 2:0:3:7 sdh 8:112 undef ready running
  |- 2:0:1:7 sdr 65:16 undef ready running
  |- 1:0:0:7 sdab 65:176 undef ready running
  `-- 1:0:3:7 sdal 66:80 undef ready running
[root@Oracle-2 ~]#

```

At this point, you can use the device labels returned by multipath to create aliases if desired. As discussed later, remapping a thin clone will not change its initial device signature, so aliases may be created once and used indefinitely for the same thin clone, regardless of what snapshot it is mapped to at any time.

6. Attaching Cloned Databases to Oracle

Now that the devices are recognized by the host, you are ready to attach the database to Oracle. In this example, the 4 LUNs comprise an Oracle ASM diskgroup called +PERF. We must first mount the diskgroup in ASM.

```

oracle@oracle-2:~
[oracle@oracle-2 ~]$ su - oracle
[oracle@oracle-2 ~]$ . oraenv
ORACLE_SID = [oracle] ? +ASM
The Oracle base has been set to /app/oracle
[oracle@oracle-2 ~]$ sqlplus / as sysasm

SQL*Plus: Release 12.1.0.2.0 Production on Wed Jan 20 15:08:05 2016

Copyright (c) 1982, 2014, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Automatic Storage Management option

SQL> alter diskgroup perf mount;

Diskgroup altered.

SQL> exit
Disconnected from Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Automatic Storage Management option
[oracle@oracle-2 ~]$ █

```

Create a new initialization parameter file for the cloned database, and adjust settings as needed to start the instance on the new host, e.g. any setting for LOCAL_LISTENER. You will likely want to create a parameter file from the source database spfile and copy it over to the new host. Also add an entry in /etc/oratab for the new database with the appropriate Oracle Home.

```

[oracle@Oracle-1 ~]$ pwd
/home/oracle
[oracle@Oracle-1 ~]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.4.0 Production on Wed Jan 20 15:13:50 2016

Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Automatic Storage Management, OLAP, Data Mining
and Real Application Testing options

SQL> create pfile = '/home/oracle/initperf.ora' from spfile;

File created.

SQL> exit
Disconnected from Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Automatic Storage Management, OLAP, Data Mining
and Real Application Testing options
[oracle@Oracle-1 ~]$ █

```

Create any necessary directories for database logging, such as \$ORACLE_BASE/admin/perf/adump and dpdump. Start the database using the newly created parameter file, and it will perform instance recovery (in the case of a crash-consistent snapshot) or will fail to open, saying it requires media recovery or that the backup must be ended (in the case of the application-consistent snapshot).


```

oracle@oracle-2:/app/oracle/product/11.2.0/dbhome_1/dbs
[oracle@oracle-2 dbs]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.4.0 Production on Wed Jan 20 15:19:45 2016

Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to an idle instance.

SQL> startup pfile='/app/oracle/product/11.2.0/dbhome_1/dbs/initperf.ora';
ORACLE instance started.

Total System Global Area 3875078144 bytes
Fixed Size 2259040 bytes
Variable Size 2281703328 bytes
Database Buffers 1577058304 bytes
Redo Buffers 14057472 bytes
Database mounted.
Database opened.
SQL> █

```

If the database does not open, issue either of the following commands via SQL*Plus or your favorite administrative tool:

```

SQL> startup pfile='/app/oracle/product/11.2.0/dbhome_1/dbs/initperf.ora'
ORACLE instance started.

```

```

Total System Global Area 3875078144 bytes
Fixed Size 2259040 bytes
Variable Size 2281703328 bytes
Database Buffers 1577058304 bytes
Redo Buffers 14057472 bytes
Database mounted.

```

```

ORA-10873: file 1 needs to be either taken out of backup mode or media recovered
ORA-01110: data file 1: '+PERF/perf/datafile/system.264.899799889'

```

```

SQL> recover database;
Media recovery complete.
SQL> alter database open;

```

```

Database altered.

```

```

SQL> exit

```

or

```

SQL> alter database end backup;

```

```

Database altered.

```



```
SQL> alter database open;
```

```
Database altered.
```

```
SQL> exit
```

7. Remap a Cloned Database

It is possible to change the state of the data inside of a thin clone to another point-in-time (snapshot). This is called remapping. The most common uses for this would be to either revert a change that was made (remap to a former snapshot) or to advance the thin clone to a more recent data set (remap to a newer snapshot).

7.1. Prepare the Database and Volumes

Similar to application consistent snapshots versus crash consistent snapshots, in order to ensure that there is no data corruption when making changes at the LUN or snapshot level, it is a best practice to offline the database and the volumes on the host.

This makes sure that nothing is in flight or in memory when the change is made. In this case, be sure to shutdown the database and dismount the diskgroup from ASM prior to remapping the thin clone.

```

oracle@oracle-2:/app/oracle/product/11.2.0/dbhome_1/dbs
[oracle@oracle-2 dbs]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.4.0 Production on Wed Jan 20 15:33:31 2016

Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Automatic Storage Management, OLAP, Data Mining
and Real Application Testing options

SQL> shutdown abort;
ORACLE instance shut down.
SQL> exit
Disconnected from Oracle Database 11g Enterprise Edition Release 11.2.0.4.0 - 64bit Production
With the Partitioning, Automatic Storage Management, OLAP, Data Mining
and Real Application Testing options
[oracle@oracle-2 dbs]$ . oraenv
ORACLE_SID = [perf] ? +ASM
The Oracle base remains unchanged with value /app/oracle
[oracle@oracle-2 dbs]$ sqlplus / as sysasm

SQL*Plus: Release 12.1.0.2.0 Production on Wed Jan 20 15:33:45 2016

Copyright (c) 1982, 2014, Oracle. All rights reserved.

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Automatic Storage Management option

SQL> alter diskgroup perf dismount;

Diskgroup altered.


SQL> █

```

7.2. Remap the Clone

Important note: Prior to remapping the contents of any clone to a different snapshot of the source LUN(s), unmount the drive from any client side applications. Failing to stop the application or unmount the LUN on the client side will have undefined consequences.

In Symphony, on the LUNs tab, select one of the LUNs associated with the snapshot group. This will change the tool bar and

introduce the clone button . Click on it and select the “Remap Thin Clone” option.

LUNs		SAN CLIENTS	SNAPSHOT RESOURCES	SNAPSHOT POLICIES	SNAPSHOTS	REPLICATIONS	SNAPSHOT GROUPS	STORAGE POOLS
LUN	LUN Type	Prod-Array-2-mg-a	Storage-Pool	Snapshot Group	Tags	Status	Total Size	Allocated Size
Oracle-500G-Thick-1	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB
Oracle-500G-Thick-1_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB
Oracle-500G-Thick-2	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB
Oracle-500G-Thick-2_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB
Oracle-500G-Thick-3	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB
Oracle-500G-Thick-3_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB
Oracle-500G-Thick-4	Thick	Prod-Array-2-mg-a	Storage-Pool	Oracle-Perf-Group		Online	500 GiB	500 GiB
Oracle-500G-Thick-4_20160120140000	Thin Clone	Prod-Array-2-mg-a	Storage-Pool			Online	500 GiB	5 GiB

Choose the desired new snapshot in the “Snapshot At:” drop down, and click on Remap at the bottom.

REMAP THIN CLONE

Current Thin Clone At: 1/20/16 2:00:00 PM

Remap for: Snapshot Group LUN

Group: Oracle-Perf-Group

Snapshot At: 1/20/16 3:00:00 PM - Scheduled_20160120150000 *

Thin Clone Name Format: {n}_{t} ? *

Thin Clone Names: Oracle-500G-Thick-2_20160120150000 Oracle-500G-Thick-4_20160120150000
Oracle-500G-Thick-1_20160120150000 Oracle-500G-Thick-3_20160120150000

Cancel
Remap

7.3. Online the Diskgroup and Database

Now that the source LUNs have been remapped to a new point-in-time (snapshot), you can simply remount the diskgroup in ASM and start the database as of the time of the snapshot. No need to rescan for the devices or even exit SQL*Plus on the ASM instance in between.

```
oracle@oracle-2:/app/oracle/product/11.2.0/dbhome_1/dbs
Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Automatic Storage Management option

SQL> alter diskgroup perf dismount;

Diskgroup altered.

SQL> alter diskgroup perf mount;

Diskgroup altered.

SQL> exit
Disconnected from Oracle Database 12c Enterprise Edition Release 12.1.0.2.0 - 64bit Production
With the Automatic Storage Management option
[oracle@oracle-2 dbs]$ . oraenv
ORACLE_SID = [+ASM] ? perf
The Oracle base remains unchanged with value /app/oracle
[oracle@oracle-2 dbs]$ sqlplus / as sysdba

SQL*Plus: Release 11.2.0.4.0 Production on Wed Jan 20 15:45:06 2016

Copyright (c) 1982, 2013, Oracle. All rights reserved.

Connected to an idle instance.

SQL> startup;
ORACLE instance started.

Total System Global Area 3875078144 bytes
Fixed Size 2259040 bytes
Variable Size 2281703328 bytes
Database Buffers 1577058304 bytes
Redo Buffers 14057472 bytes
Database mounted.
Database opened.
SQL>
```