

ORACLE®

Rapid database cloning using SMU and ZFS Storage Appliance

How Exalogic tooling can help

Jacco H. Landlust
Platform Architect Director
Oracle Consulting NL, Core Technology
December, 2014

Presented at



Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Program Agenda

- 1 Overview
- 2 Features & Terminology
- 3 Use Case: Implementing SMU with Exadata and ZFS
- 4 Demo
- 5 Q & A

Overview



- **Engineered to work together**
 - Sun ZFS Storage Appliance - Oracle Database integration
 - ZFS Storage snapshot, clone, rollback capabilities
- **Database Storage Efficiency**
 - Rapid and Efficient backups, restores and provisioning
 - Empowers secondary processing such as Dev&Test, Reporting
- **Simplifies Database Storage Management**
 - Browser User Interface and CLI affords rapid familiarity
 - Setup DR environment with optional Remote Replication

Comparing with Snap Clone from Enterprise Manager

- **Different License**
 - Snap Management Utility versus Cloud Management Pack
- **SMU is stand alone, resulting in differences impacting**
 - Clustering
 - Security
 - Auditing
 - Logging
- **SMU is specific for ZFS whereas Snap Clone from EM is storage agnostic**

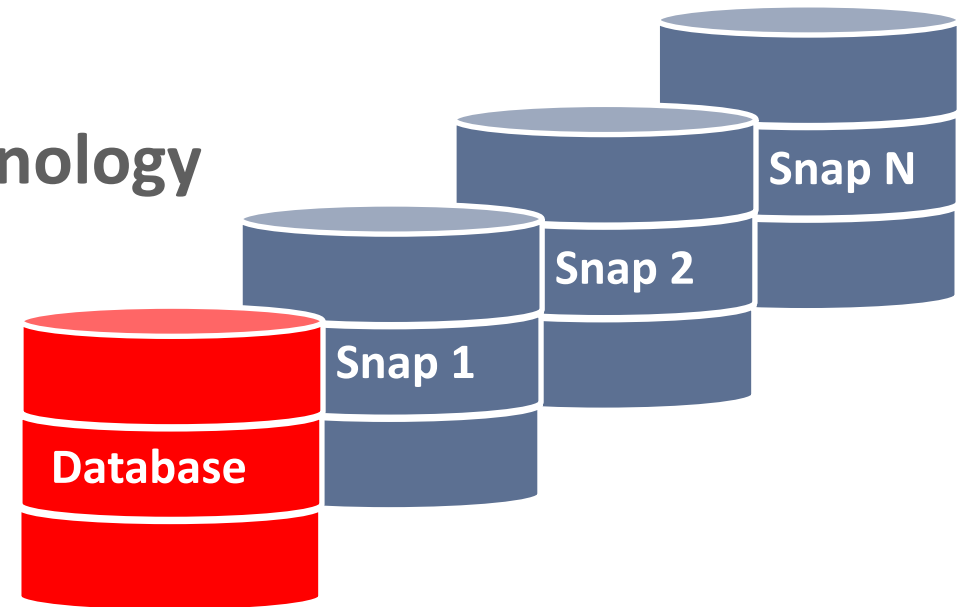


Setting the Scene

Features & Terminology

Database Snapshot Backups

- **Hot: Online**
 - No interruption to the database
- **Cold: Offline**
 - Database shutdown
- **Sun ZFS Storage Appliance Snapshot technology**
 - Unlimited snapshots
- **Rapid Familiarity**
 - Browser User Interface for ease of use
 - Powerful Command Line Interface



Database Types

- **Unified Storage Modality**

- File Based Storage:

- Network Mounted Shares
- NFS & dNFS supported

- Block level Storage

- Automatic Storage Manager - iSCSI
- Utilizing DB technology


- **Database Deployment Modality**

- Single Instance
- RAC (cluster)



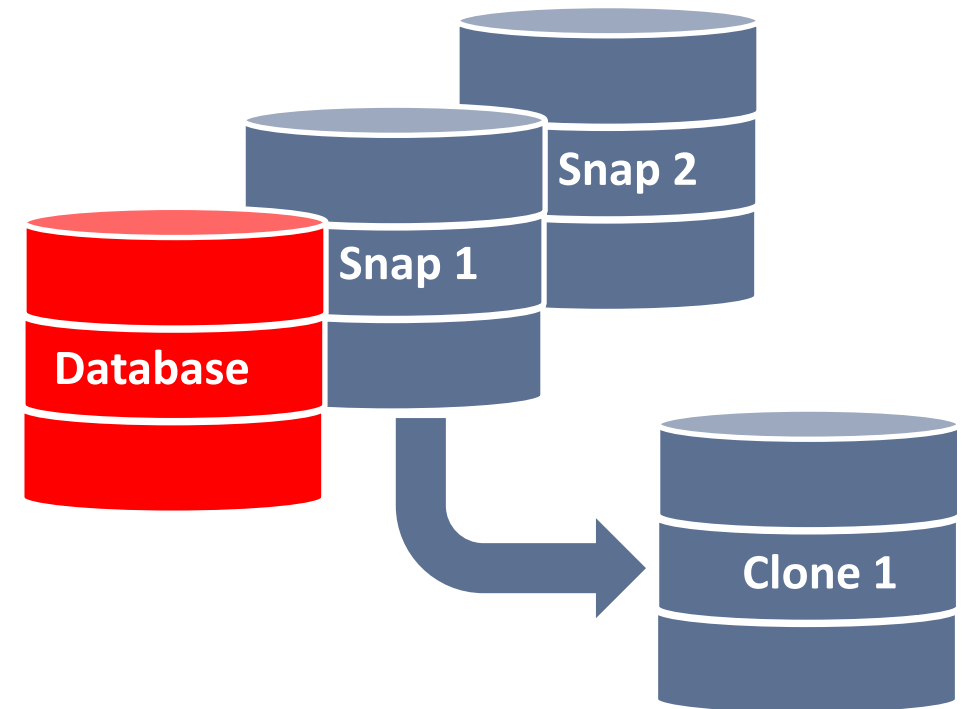
Scheduler & Policy Management

- **Database Storage Automation**
 - Automate snapshots
 - Schedule snapshots
 - Flexible snapshot Intervals
 - Minutes, Hours, Days, Weeks, Months
 - Set retention policies
 - Automate deletion of snapshots
 - Ease Snapshot Management
 - Enforces Snapshot Lifecycle policies
 - Alert Notification (Email)



Snapshot Operations

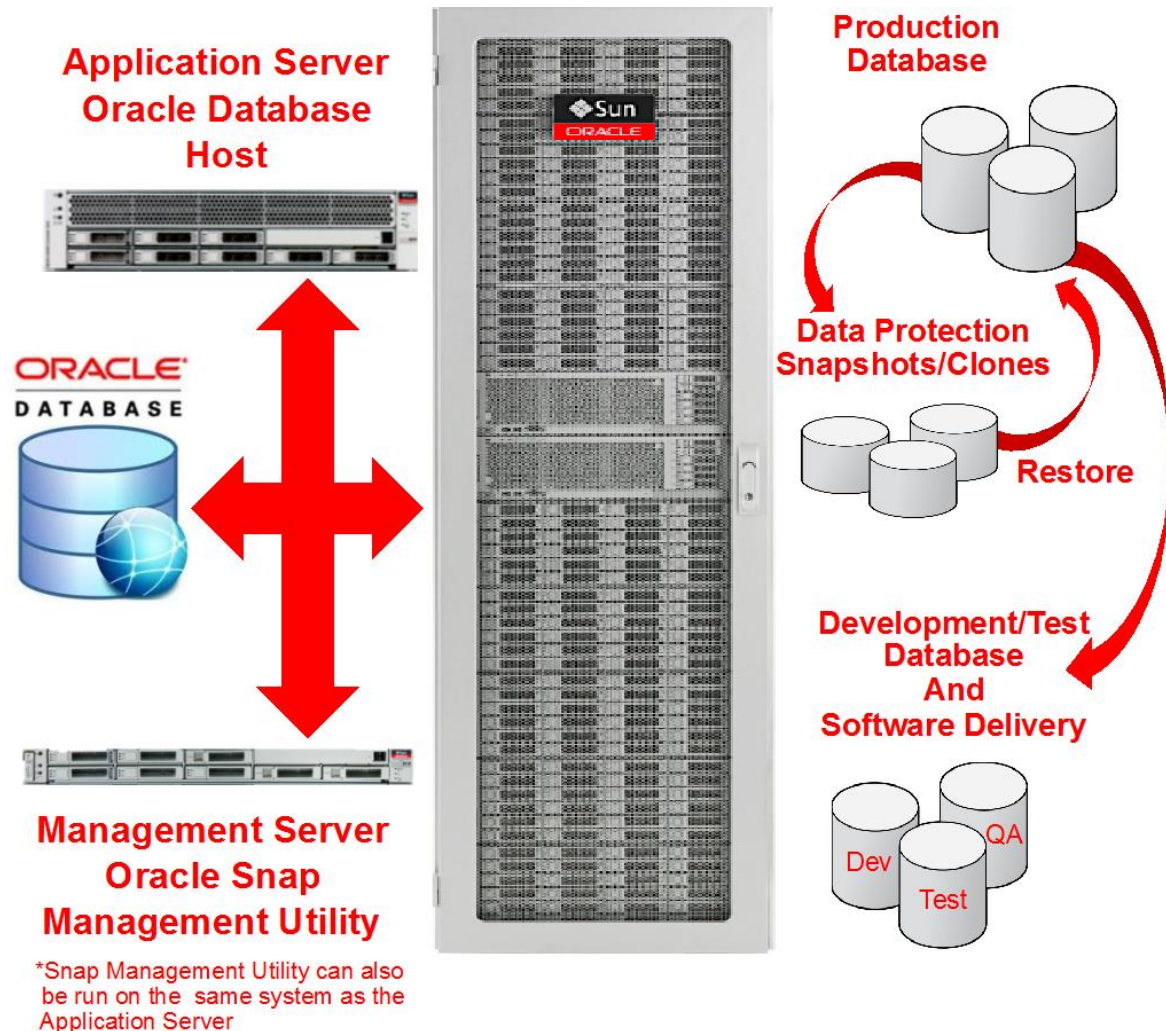
- **Snapshot/Clone Management**
 - Manage Snapshot/Clone attributes
 - Rename snapshots
 - Delete snapshots
 - De-provision Clone
- **Database Management**
 - Restore database
 - Clone databases
- **Heterogeneous OS Environments**
 - Windows, Solaris, Linux



Terminology

- **ZFS Snapshot Backup**
 - Read-Only Point in Time Copy
- **ZFS Clone**
 - Read-Write Point in Time Snapshot Copy
- **ZFS Restore**
 - Restored to any 'point in time' snapshot of:
 - File based modality – Filesystem
 - Block level modality – LUN
 - Storage Efficient
 - Deletes all subsequent Snapshots
- **ZFS Replication**
 - Asynchronous snapshot data
 - Local Copy – Current Storage Host
 - Remote Copy – Remote Storage Host*
 - Bandwidth Efficient:
 - Only transfer/updates changed blocks
- **ZFS Database Snapshot/Clone**
 - Snapshot based backup/clone of Oracle Database

*Remote host must be a Sun ZFS Storage Appliance



Limitations

- All database editions except Express Edition
- RAC One Node is not supported
- External redundancy diskgroups for ASM storage types
- Each database must use its own separate set of shares
- Database shares must reside within a single storage head
- ASM cloning requires Oracle Database 11.2.0.2.0 or later
- Each database must have certain file layouts that depend on the type of snap backups that will be taken

CRISIS COUNSELING

**THERE IS HOPE
MAKE THE CALL**

**THE CONSEQUENCES OF
JUMPING FROM THIS
BRIDGE ARE FATAL
AND TRAGIC.**

Use Case

Implementing SMU with Exadata and ZFS

Generic Use Cases

- **Snapshot based backup/restore of Oracle Databases**
 - hot, cold
- **Clones of (production) databases**
 - Running reports, development, QA, etc.
- **Heterogeneous environments**
 - Migrate clones to other hardware, OS's
- **DB Clone from RMAN Image Copy**



Use Case at Customer: Clone(s) of databases for Reporting

- **Application description**

- risk management on processing and back trading

- **Business Requirement**

- (re-)run checks on end of days

- Investigate findings on trades from previous days (up to 10 days back)

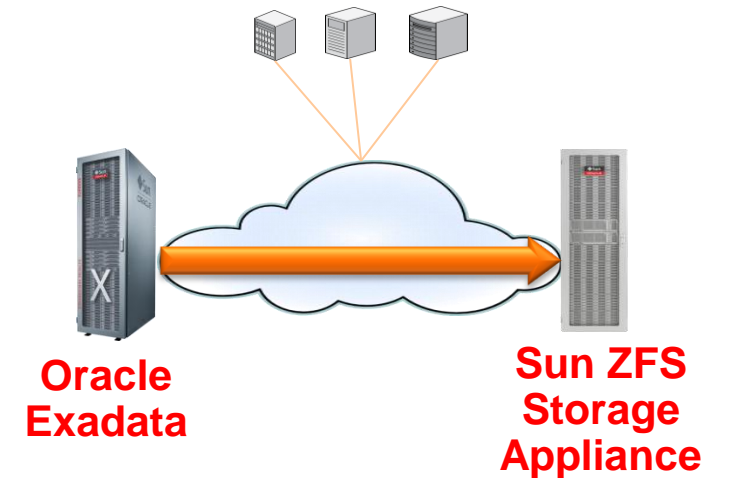
- **Database Details**

- Approximately 4 TB

- 90 to 95 % of the records are changed on daily basis

Environment at Customer in Numbers

- ZS3-2 connected via Infiniband to Exadata
- 10Gbit to Core Network
- 256GB DRAM Cache, 1.6TB SSD Read Cache
- 80 x 4 TB HC Disks + 4 x Logzilla



Type	NSPF	Width	Spares	Data Drives	Raw (TB)	Usable (TiB)
Mirror	True	2	4	76	152	136
Raidz1	False	4	4	76	228	204

Implementation Plan: Use regular RMAN backup from Production

- Does not interfere with regular processes
- Requires image copy backup
- Data files on share 1
- Archives and control file on share 2
- Only **ONE** full backup on the shares

RMAN Performance

- **RMAN Image Backup**
 - 150MB/s per max per RMAN channel (Bigfile)
 - 1.15GB/s per DB node with 16 channels (9 Active)
 - 4 IB connections to ZFS
 - ZFS 1MB recordsize
 - ZFS in Throughput mode
- **RMAN Duplicate from active**
 - 115MB/s to ZFS (no MTU tuning)
 - 256MB/s to ASM

RMAN Performance

- 2 – 4 RMAN Channels per ZFS Disktray to get best performance
- Image Backup can not split data files so 8 – 16 data files needed to fill channels
- ZFS record-size does not seem to matter much
- Number of IB connections has not got a big impact as a single connection is not saturated

Other Caveats

- **Timed starting**
 - Only snap clone after a specific transaction is hard to implement
- **No control over init.ora**
 - RMAN Restore resulted in a new (default) init.ora
 - 2GB SGA and AMM configured

Result

Failure

Alternative plan: Setup cascading stand-by on ZFSSA

- Only snapshot and clone required
- More granular snapshots possible
- Slightly more complicated setup;
no data guard knowledge at customer present
- Downside: possible impact on production process (risk: low)




Demo

ZFS Storage Appliance Simulator

- Virtualbox Appliance
- Configured with 2048 MB of memory and 125 GB of dynamically allocated disk space
- See <http://www.oracle.com/us/dm/h2fy11/simulator5-minute-guide-192152.pdf> for how to setup the ZFSSA
- Sun ZFS Storage Appliances must be running Appliance Kit (AK) firmware version 2011.1.5 or later for SMU

“Application Server”

- Runs Oracle Linux 6.5
- 6GB RAM
- Oracle Database 11.2.0.4 Enterprise Edition



Thanks for all the fish

Questions?

Oracle Snap Management Utility for Oracle Database Resources

Oracle Snap Management Utility for Oracle Database

<http://www.oracle.com/us/products/servers-storage/storage/nas/snap/overview/index.html>

Sun ZFS Storage Appliance

<http://www.oracle.com/us/products/servers-storage/storage/nas/overview/index.html>

Oracle Database

<http://www.oracle.com/us/products/database/overview/index.html>

Oracle Engineered Systems

<http://www.oracle.com/us/products/engineered-systems/index.html>

Snap Management Utility for the Oracle Database - Information and Troubleshooting (Doc ID 1522925.1)

Hardware and Software Engineered to Work Together

ORACLE®