

Oracle Maximum
Availability Architecture

An Oracle White Paper
July 2012

Oracle Exadata Database Machine Consolidation: Segregating Databases and Roles

Introduction	2
Goal.....	2
Scope	2
Best Practices Summary	4
Application Database Isolation Objectives	5
Separate Installation Owners.....	6
Job Role Separation	6
Diskgroup Isolation	7
Application Database Isolation Configuration.....	8
Initial Exadata Database Machine Installation.....	9
Operating System Level Configuration	9
Create the Griddisks and ASM Diskgroups.....	11
Configure Exadata Database-scoped Security	13
Further Considerations	13
Conclusion	13
Appendix	14
Test Environment Details.....	14
Operating System Group and Account Configuration Examples ..	15
Griddisk and ASM Diskgroup Creation	17
Configuring Exadata database-scoped Security Example.....	21
dbScoped_EBS.sh Generated Script.....	28
References	31
Change Record	31

Introduction

This paper is focused on the aspects of segregating databases from each other in a platform consolidation environment on an Oracle Exadata Database Machine. Platform consolidation is the consolidation of multiple databases on to a single Oracle Exadata Database Machine. When multiple databases are consolidated on a single Database Machine, it may be necessary to isolate certain database components or functions in order to meet business requirements and provide best practices for a secure consolidation. In this paper we outline the use of Oracle Exadata database-scoped security to securely separate database management and provide a detailed case study that illustrates the best practices.

Goal

The goal of this case study example is to isolate and secure each of the application databases on Exadata Database Machine so that the grid infrastructure and each application database are administered separately.

Scope

There are many other aspects to implementing a platform consolidation that are outside the scope of this paper. Some of those aspects are capacity planning, workload management, improved manageability, lower total cost of ownership (TCO), more efficient utilization of computing resources, improved service levels and availability, more effective skill utilization, and reduced need for floor space in corporate and divisional data centers. Many of these topics will be covered in separate papers. Continue to check the [MAA Best Practices - Exadata Database Machine](#) site for any updates or new papers.

This paper documents the case where each application database is administered by separate database administrators. If multiple application databases are administered by the same database administrator(s) then there may be reason to group those databases under the same operating system account, but this case is not covered in this paper. Even in the case of a database administrator being responsible for multiple databases it may still make sense to setup separate administration as detailed in this paper.

This example uses Linux 64-bit, however the same concepts can be applied by administrators using the Solaris Operating System. The following application workloads

were used in this example, running on their own separate Oracle Database 11g release 2 (11.2.0.2) databases:

1. Oracle E-Business Suite
2. Peoplesoft Enterprise Human Capital Management
3. Siebel Customer Relationship Management (CRM)

The application database autonomy will be achieved through the use of job role separation through separate operating system accounts, operating system groups for Oracle ASM users and using Oracle Exadata database-scoped security. An overview of the planned deployment is shown in Figure 1.




Figure 1 Case Study Deployment Overview

Oracle Exadata Database machine implements basic operating system hardening policies that are described in the “Oracle Exadata Storage Server Software User’s Guide”

[5]. Oracle Exadata Database Machine should be implemented within your existing security architecture.

This paper does not delve into Oracle database security itself. Oracle provides a defense-in-depth security architecture that provides an “inside-out “security strategy covering both protective and detective security controls that can be deployed in combination or individually. Oracle’s defense-in-depth approach to security is critical to safeguarding data and is a key component of an effective and compliant security strategy. Both Oracle Database Vault and Oracle Advanced Security can be deployed in the Oracle Exadata Database Machine in the same way they are deployed on any other Oracle Database configuration (single instance, Oracle Real Application Clusters (RAC) or Oracle Data Guard environments). Use of the Oracle defense-in-depth security architecture features is beyond the scope of this paper but should be reviewed and understood by reading the “[Cost Effective Security and Compliance with Oracle Database 11g Release 2](#)” Oracle white paper [6]. Understanding the Oracle defense-in-depth security architecture is critical to securely deploying a consolidated environment.

Best Practices Summary

This section summarizes the best practices derived from this case study to separate roles and database access for each application database. Note that this case study documents the details for full isolation of each application database and that this should meld into an existing security architecture.

There may be varying requirements where not all of these best practices are required. For example, some databases are under the same administrator, under the same operating system account and share the same Oracle database software installation. These cases are not explicitly covered but you should be able to use some or all of these best practices to configure your environment to meet your requirements.

The details of implementing each best practice are found in the subsequent sections.

- To separate Grid/ASM administration from database administration, request a "job role-separated environment" [3] for the initial Exadata installation.

This facilitates separating the Oracle Grid Infrastructure administrative privileges from Oracle Database administrative privileges. There is a single grid infrastructure installation for the database machine. See "[Initial Exadata Database Machine Installation](#)" for details.

- To separate database administration for different databases, create a separate operating system user account, operating system group and database ORACLE_HOME for each database.

See "[Operating System Level Configuration](#)" for details.

- To separate database storage create separate diskgroups

Separate data and fast recovery area (FRA) disk groups can be created for each database or group of databases. It is best to define the diskgroups to span all cells and disks which will provide a balanced configuration and offer the highest levels of performance and availability for each diskgroup.

By having its own set of griddisks and pair of diskgroups for the data and FRA diskgroups, isolation for each application database is provided. For details see "[Create the Griddisks and ASM Diskgroups](#)".

- To control administrator access to diskgroups, use Exadata database-scoped security

By implementing Exadata database-scoped security it is possible to forbid database administrators access to diskgroups that are not used for their databases. Oracle Exadata database-scoped security controls which databases can access specific grid disks that compose Oracle ASM disk groups, while still sharing all of the physical I/O resources. See "[Configure Exadata Database-scoped Security](#)" for details.

- Use the Oracle defense-in-depth security architecture features.

While beyond the scope of this paper, you should review and understand Oracle defense-in-depth security by reading the "[Cost Effective Security and Compliance with Oracle Database 11g Release 2](#)" paper and apply accordingly to your databases to implement a secure database.

Application Database Isolation Objectives

To reiterate, the goal in this case study is secure application database isolation on an Oracle Exadata Database Machine. The primary requirement is to administer the grid infrastructure and each application database separately and ensure one administrator cannot access another administrator's software or database. Thus, in addition to the grid infrastructure having its own

operating system account and ORACLE_HOME, each application database can remain autonomous from any other application database by having their own of each of the following:

- Operating system account
- Oracle database software
- Oracle Automated Storage Management (ASM) diskgroups

This will be accomplished through the following tactics:

- Separate Installation Owners
- Operating system job role separation
- Oracle ASM Diskgroup isolation using Exadata database-scoped security

A short overview of these objectives is given below and then a detailed configuration example follows.

Separate Installation Owners

In addition to separating the Oracle Grid Infrastructure installed under its own operating system account, each of the application databases will be setup under separate operating system user accounts so that each of the application database installations has a different owner. This facilitates isolation of the ORACLE_HOME's and configuring Exadata database-scoped security to isolate the Oracle ASM diskgroups from each other.

Job Role Separation

To further facilitate separate grid infrastructure (ASM administration primarily) and database administration roles, it is necessary to create operating system groups that map to the different roles. A job role separation privileges configuration facilitates separate roles for managing and using Oracle ASM. This configuration of groups and users divides administrative access privileges to the Oracle ASM installation from other administrative privileges users and groups associated with the application database Oracle installations. Administrative privileges access is granted by membership in separate operating system groups, and installation privileges are granted by using different installation owners for each Oracle installation. This is all detailed in "[Oracle® Grid Infrastructure Installation Guide 11g Release 2 \(11.2\) for Linux – 2.5 Creating Groups, Users and Paths for Oracle Grid Infrastructure](#)" [3].

Additionally, separate operating system groups will also be used for Oracle ASM users. This is detailed in "[Oracle® Automatic Storage Management Administrator's Guide 11g Release 2 \(11.2\) – Chapter 3, Using Separate Operating System Groups for Oracle ASM Users](#)" [4].

Diskgroup Isolation

To isolate database administration of each application database and facilitate the use of database-scoped security it is necessary to allocate each application database its own pair of diskgroups for data and the fast recovery area (FRA). Allocating diskgroups to each application database for data and FRA disk groups requires proper space requirements planning. Also check the [Oracle Automatic Storage Management Administrator's Guide \[4\]](#) for ASM storage limits. Currently Oracle ASM has a limit of 63 disk groups in a storage system.

Like most consolidation scenarios, this example used existing applications so it was relatively simple to calculate the space requirements. This paper will not go into the details of the capacity planning. Once the space requirements are understood for each application then you can derive the griddisk size based on the number of disks available. As stated earlier, each application will have its own pair of data and FRA Oracle ASM diskgroups which also implies a similar set of griddisks for each Oracle ASM diskgroup on each Exadata storage cell. **To provide a high level of I/O performance, diskgroups will be configured using equal sized disk stripes (grid disks) within a diskgroup on every disk on every Exadata cell.** The DATA diskgroups will be configured with high redundancy to provide maximum protection from cell outages and storage failures (see [MAA Best Practices for Oracle Exadata Database Machine \(technical white paper\) \[11\]](#) for details). Note that online redo logs and standby redo logs should go on the high redundancy diskgroup and not be multiplexed. The FRA will be configured with normal redundancy. The order of creation will be:

1. DATA diskgroup griddisks
2. Space for future growth or new databases
3. FRA diskgroup griddisks

This order of creation is important to note since the griddisks created first (with a lower offset) will be on the outer tracks of the physical disks and consequently get slightly better performance. You can also specify the OFFSET as part of the griddisk create commands and then the order is not pertinent as you can place the griddisks at a specific offset. This is depicted below in Figure 2, which shows an example of the quarter rack, three Exadata Storage cells, used in this case study. The details of creating this configuration are included in the appendix at "[Griddisk and ASM Diskgroup Creation](#)".




Figure 2 Consolidation Griddisk / Diskgroup Layout

Oracle Exadata database-scoped security [5] (see chapter 4) can then be configured to isolate the ASM diskgroups from each other so that only the authorized DBA has access to the application’s disk groups. Database-scoped security configures access to specific grid disks on cells for specific database clients of an Oracle ASM cluster. This security mode assists you to control which databases can access specific grid disks that compose Oracle ASM disk groups. Setting up Oracle database-scoped security first requires configuring ASM-scoped security for your initial security mode. Then database-scoped security for specific database clients and grid disks can be configured. After setting up database-scoped security among the database clients and grid disks, only those specific grid disks are available to the specified database clients.

Application Database Isolation Configuration

This section will go through the steps to configure the environment for secure application database isolation on the Oracle Exadata Database Machine. The steps will entail:

- Preparing for the [initial Exadata Database Machine installation](#)
- Post initial installation [operating system configuration](#)
- [Create the Exadata griddisks and ASM diskgroups](#)

- [Configure Exadata Database-scoped Security](#)
- Review [further security requirements](#) and configuration

Initial Exadata Database Machine Installation

When preparing for the initial installation be aware of the following:

- This case study requires a separate operating system user for the grid infrastructure installation. Ensure the requirement for a job role-separated environment is communicated to the install team prior to the initial machine setup since the initial machine setup process of using the OneCommand utility [10] will then create the appropriate operating system accounts and groups as detailed at [Creating Groups, Users and Paths for Oracle Grid Infrastructure](#). Additional application specific operating system users and groups will be created later.
- The initial installation will configure a pair of DATA and FRA diskgroups. These will be dropped when configuring application database diskgroup pairs.

NOTE

If this is for an existing installation that you now wish to separate roles for then you can leave the current grid infrastructure as is and move the application database ORACLE_HOMEs to a new operating system account.

Operating System Level Configuration

To isolate the application databases, the separate job role install requires creating additional operating system groups and users in addition to the standard job role separated install.

- Reference documentation [Creating Groups, Users and Paths for Oracle Grid Infrastructure](#) and discuss the OS groups and their roles

Operating System Groups

In our case the operating system group configuration will be as follows:

OS Group	Role	Database Role
Initial Setup		
oinstall	Install	Required for installing ORACLE_HOME's. You must be a member of this group to install or patch the Oracle software.
dba	Generic oracle OS user database administration	Typically will be used to install the Enterprise Manager agent. Separate application database DBA OS groups are created post-install.
asmadmin	OSASM	This is a required group. Create this group as a separate group if you want to have separate administration privilege groups for Oracle ASM and Oracle Database administrators. Members of the OSASM group can use SQL to connect to an Oracle ASM instance as SYSASM using operating system authentication. The SYSASM privileges permit mounting and dismounting disk groups, and other storage administration tasks. SYSASM privileges provide no access privileges on an RDBMS instance.

OS Group	Role	Database Role
asmdba	OSDBA for ASM	Members of the ASM Database Administrator group (OSDBA for ASM) are granted read and write access to files managed by Oracle ASM. The Oracle Grid Infrastructure installation owner and all Oracle Database software owners must be a member of this group, and all users with OSDBA membership on databases that have access to the files managed by Oracle ASM must be members of the OSDBA group for ASM.
asmoper		This is an optional group. Create this group if you want a separate group of operating system users to have a limited set of Oracle ASM instance administrative privileges (the SYSOPER for ASM privilege), including starting up and stopping the Oracle ASM instance. By default, members of the OSASM group also have all privileges granted by the SYSOPER for ASM privilege.
Post-initial Setup required		
dba_ebs	OSDBA for EBS	SYSDBA for EBS database
dba_psft	OSDBA for PeopleSoft	SYSDBA for PeopleSoft database
dba_siebel	OSDBA for Siebel	SYSDBA for Siebel database

TABLE 1 OPERATING SYSTEM GROUPS

For the post-initial setup see [“Create Post-install OS Groups”](#) for a detailed example.

Operating System Accounts

The operating system accounts are configured partly by the initial install done with the OneCommand utility [10] and partly post-initial-install. For the application database administration accounts, direct access will be restricted to using the sudo feature from specified user accounts. Thus, application database administration accounts may only be accessed via sudo to create an audit trail of their usage. A sudo setup example is included at [“Setup sudo Access”](#).

For this example the following operating system account configuration was set up:

OS User	Role	OS Group(s)	Sudo Access	Notes
root	System Administrator	0 (root), 1 (bin), 2 (daemon), 3 (sys), 4 (adm), 6 (disk), 10 (wheel)		
Setup by initial install				
oracle	To manage the EM agent only	oinstall dba racoper asmdba		
grid	Oracle ASM administrator / Oracle Grid Infrastructure home	oinstall asmdba asmoper asmadmin		Owns the GI stack, including the ASM instances
Post- install				
oracle_ebs	EBS DB Software Owner and DBA	dba_ebs, oinstall, asmdba	user1	No login, only sudo access
oracle_psft	PeopleSoft DB Software Owner and DBA	dba_psft, oinstall, asmdba	user2	No login, only sudo access
oracle_siebel	Siebel DB Software Owner and DBA	dba_siebel, oinstall, asmdba	user3	No login, only sudo access
user1			oracle_ebs	Login and sudo to work
user2			oracle_psft	Login and sudo to work
user3			oracle_siebel	Login and sudo to work

TABLE 2 OPERATING SYSTEM ACCOUNTS

For the post-initial setup see “[Create Post-install OS Accounts](#)” for a detailed example. The example in the appendix includes the necessary setup for shell limits in the /etc/security/limits.conf file and configuring sudo access for the application accounts.

Lastly, ensure that you review the generic DB security best practice paper, “[Oracle Database Security Checklist](#) - Technical Whitepaper” for further security best practices.

Create the Griddisks and ASM Diskgroups

When configuring the Oracle ASM diskgroup pairs to support the application isolation requirement the following was established:

- Each application has its own pair of Oracle ASM disk groups, DATA and FRA
- The storage server grid disk configuration should span all cells to maximize the I/O bandwidth

- The placement (griddisk offset) is important because the griddisks created on the lower griddisk offset will be on the outer portion of the physical disk and therefore have better performance.

With that in mind then these questions should be asked:

- How much space is available to divide up?
- What are the application database disk space requirements?
 - Disk space per application (including growth rates)
 - Oracle ASM high redundancy versus normal redundancy
 - How to fit the configuration within the security model?

In this Exadata quarter rack environment there are three Exadata Storage Servers with 600 GB SAS high performance disks. That equates to 36 disks with about 21 TB of raw storage. In this example high redundancy is used for the data disk groups and normal redundancy for the fast recovery area disk groups. The application raw disk space requirements included a 2% growth factor and additional space for other consolidation and are listed in Table 3. The raw disk space requirements take into account the high redundancy for DATA and normal redundancy for FRA by appropriately multiplying the actual requirement by 3 and 2, respectively. With these requirements it was determined that an Exadata quarter rack had sufficient disk space. The details of capacity planning and calculating the space requirements in detail are beyond the scope of this paper.

Disk Group (Raw Sizes in GB)	EBS	Peoplesoft	Siebel	Additional Consolidation	TOTAL
DATA	1,988	3,290	1,200	1,467	7,945
FRA	3,226	2,581	2,128	2,641	10,576
TOTAL	5,214	5,871	3,328	4,108	18,521

TABLE 3 APPLICATION RAW DISK SPACE REQUIREMENTS

Dividing each of those numbers by the number of disks, 36, will then give the required size for each griddisk. It should also be noted that the initial disk group created for the Oracle RAC Oracle Cluster Registry (OCR) and voting disks is left as is from the initial system install.

Using the raw space requirements and adding 1 GB to each result gives the griddisk sizes shown in Table 4. Using those numbers the griddisks are created and then the ASM disk groups are

created as detailed in the appendix under “[Griddisk and ASM Diskgroup Creation](#).” Note that the “Additional Consolidation” cell disk space was allocated as griddisks between the DATA and FRA and no associated ASM disk groups were created at this point and left for future considerations.

Griddisks (Sizes in GB)	EBS	Peoplesoft	Siebel	Additional Consolidation
DATA	56	91	34	42
FRA	90	72	60	83

TABLE 4 GRIDDISK SIZE REQUIREMENTS

Configure Exadata Database-scoped Security

Exadata Cell data security is implemented by controlling which Oracle ASM clusters and database clients can access specific grid disks on storage cells. By default, all database and Oracle ASM instances have access to all storage cell grid disks. The goal of configuring Exadata Database-scoped Security is to restrict griddisk (diskgroup) access by unauthorized users. Since each application database has its own operating system account and own ORACLE_HOME this facilitates configuring database-scoped security for each application database and their associated griddisks which in turn restricts diskgroup access. See the appendix, “[Configuring Exadata database-scoped Security Example](#)”, for the complete details of configuring database-scoped security.

Further Considerations

If you are using Oracle Enterprise Manager (EM) to manage and monitor your environment then the same roles should be administered in EM. Also see “[Support Note:1110675.1 - Monitoring Exadata Database Machine using Enterprise Manager - contains updated monitoring suggestions for each component](#)”, for details of configuring the EM Exadata plug-ins.

Also, to reiterate, the Oracle defense-in-depth security architecture should be reviewed and understood by reading the “[Cost Effective Security and Compliance with Oracle Database 11g Release 2](#)” Oracle white paper [6]. Understanding the Oracle defense-in-depth security architecture is critical to securely deploying a consolidated environment.

Conclusion

This paper covered the extreme case of segregating all databases and their administration from each other with separate diskgroups in a consolidated environment. Administration roles may not always be this clear cut and there will be hybrid implementations that include multiple databases under a single administrator or group of administrators that necessitate use of a shared

database software installation and/or shared Oracle ASM diskgroups. There also may be chargeback requirements that are part of implementing a consolidated environment and influence the configuration. These and any other variations need to be factored into implementing a robust and secure consolidated deployment that fits within your existing security architecture. This paper should provide a basis for those variations as well as for the case detailed herein. Continue to check the [MAA Best Practices - Exadata Database Machine](#) site for any updates or new papers.

Appendix

Test Environment Details

The hardware and software details for the test environment are as follows:

Target Database System

Oracle Exadata Database Machine Oracle Linux quarter rack

- Compute Nodes
 - Names: sclczdb01 and sclczdb02
 - IP Addresses (the two high-order IP address octets, xxx.xxx, are used to protect IP addresses)
 - sclczdb01 xxx.xxx.74.168
 - sclcz01-vip xxx.xxx.77.181
 - sclczdb02 xxx.xxx.74.169
 - sclczdb02-vip xxx.xxx.77.182
 - sclcz-scan xxx.xxx.77.177, xxx.xxx.77.178, xxx.xxx.77.179
 - \$ host sclcz-scan

 - sclcz-scan.us.oracle.com has address xxx.xxx.77.177
 - sclcz-scan.us.oracle.com has address xxx.xxx.77.178
 - sclcz-scan.us.oracle.com has address xxx.xxx.77.179

 - See <http://www.oracle.com/technetwork/database/clustering/overview/scan-129069.pdf> for understanding the Oracle RAC Single Client Access Name (SCAN) configuration.
- Exadata software version 11.2.2.3.2
- Grid ORACLE_HOME 11.2.0.2
 /u01/app/11.2.0/grid
- ASM ORACLE_SID=+ASM1 and +ASM2 respectively
- Oracle Linux 2.6.18-194.3.1.0.2.el5 x86_64
- Oracle Grid Infrastructure 11g Release 2 Enterprise Edition (11.2.0.2)
- Oracle Database 11g Release 2 Enterprise Edition (11.2.0.2)

- 2 Sun Fire X4170 Quad-Core Intel Xeon® E5540 Processors (2.53 GHz)
Note: this test was on an Exadata V2 system which is no longer available. See <http://www.oracle.com/exadata> for current availability.
- 72 GB memory
- Disk Controller HBA with 512MB Battery Backed Write Cache
- 4 x 146 GB SAS 10,000 RPM disks
- Dual-Port QDR Infiniband Host Channel Adapter
- 4 Embedded Gigabit Ethernet Ports
- Storage:
3 Exadata Storage Servers (sclczcel01-03), each with:
 - Exadata software version 11.2.2.3.2
 - 2 Quad-core Intel Xeon E5540 (2.53GHz) processors
Note: this test was on an Exadata V2 system which is no longer available. See <http://www.oracle.com/exadata> for current availability.
 - 12 x 600 GB High Performance SAS drives
 - 384 GB Exadata Smart Flash Cache

Operating System Group and Account Configuration Examples

All these examples are for Oracle Linux.

Create Post-install OS Groups

```
# dcli -g ~/dbs_group -l root -x setupOSgroups.sh
```

setupOSgroups.sh

```
#!/bin/ksh
#
groupadd dba_ebs -g 1010
groupadd dba_psft -g 1011
groupadd dba_siebel -g 1012
#
# Add individual groups
groupadd user1 -g 1110
groupadd user2 -g 1111
groupadd user3 -g 1112
```

Create Post-install OS Accounts

First create the two scripts and then execute them on all compute nodes.

setupOS_dbaUsers.sh

```
#!/bin/ksh
#
export PASSWD=welcome1
export USER_LIST="oracle_ebs oracle_psft oracle_siebel"
# Set the USER_ID to the next available id from /etc/passwd or an unused value
export USER_ID=1002
```



```

export LC=/etc/security/limits.conf
# Backup /etc/security/limits.conf
cp $LC ${LC}_`date +%Y-%m-%d_%H:%M:%S`
#
for username in `echo $USER_LIST`
do
    export user_suffix=`echo $username | cut -f2 -d_`
    useradd $username -G dba_$user_suffix,oinstall,asmdba -u $USER_ID -g oinstall \
-d /home/$username
    echo "$username:$PASSWD" | chpasswd
    chown -R $username:oinstall /home/$username
    # Update /etc/security/limits.conf
    echo " " >> $LC
    echo "$username          soft   core      unlimited" >> $LC
    echo "$username          soft   core      unlimited" >> $LC
    echo "$username          soft   nproc    131072" >> $LC
    echo "$username          hard   nproc    131072" >> $LC
    echo "$username          soft   nofile   131072" >> $LC
    echo "$username          hard   nofile   131072" >> $LC
    echo "$username          soft   memlock  55520682" >> $LC
    echo "$username          hard   memlock  55520682" >> $LC
    (( USER_ID+=1 ))
done

```

setupOSUsers.sh

```

#!/bin/ksh
#
export PASSWD=welcome1
export USER_LIST="user1 user2 user3"
# Set USER_ID to the next available id from /etc/passwd or an unused value
export USER_ID=1102
#
for username in `echo $USER_LIST`
do
    useradd $username -g $username -u $USER_ID -d /home/$username
    echo "$username:$PASSWD" | chpasswd
    (( USER_ID+=1 ))
done

```

Create Post-install OS Accounts

```

# dcli -g ~/dbs_group -l root -x setupOS_dbaUsers.sh
# dcli -g ~/dbs_group -l root -x setupOSUsers.sh

```

Setup sudo Access

Edit /etc/sudoers file making the following changes.

1. Set the sudo Log file specification
(by default sudo is logged to /var/log/secure)
Defaults logfile=/var/log/sudo.log
2. Setup user aliases
User_Alias EBS = user1
User_Alias PSFT = user2
User_Alias SIEBEL = user3
3. Setup su command aliases
su Aliases

```

Cmnd_Alias EBS_SU = /bin/su - oracle_ebs, /bin/su oracle_ebs
Cmnd_Alias PSFT_SU = /bin/su - oracle_psft, /bin/su oracle_psft
Cmnd_Alias SIEBEL_SU = /bin/su - oracle_siebel, /bin/su oracle_siebel

```

4. Setup user access

```

# Setup user access using User_alias settings
#EBS ALL = /bin/su - oracle_ebs, /bin/su oracle_ebs
EBS ALL = EBS_SU
#PSFT ALL = /bin/su - oracle_psft, /bin/su oracle_psft
PSFT ALL = PSFT_SU
#SIEBEL ALL = /bin/su - oracle_siebel, /bin/su oracle_siebel
SIEBEL ALL = SIEBEL_SU

```

5. Remove ssh access to the application database accounts so that they can only be accessed via sudo

- a. Edit `/etc/ssh/sshd_config` and add the following line

```
DenyUsers oracle_ebs oracle_psft oracle_siebel
```

- b. Restart the ssh daemon

```
service sshd restart
```

Griddisk and ASM Diskgroup Creation

Create Griddisks

All the following are run from a user on the management node (may be a compute node) that has established user equivalence to the celladmin user on the storage cells unless otherwise noted.

1. From the grid infrastructure OS user (grid) drop any unused diskgroups from the initial install.
 - a. As the grid infrastructure owner, list existing diskgroups

```

ASMCMD> lsdg
State   Type      Rebal  Sector  Block      AU   Total_MB  Free_MB
Req_mir_free_MB  Usable_file_MB  Offline_disks  Voting_files  Name
MOUNTED  HIGH      N          512    4096    4194304  2064384  678372
375342          101010          0          0          N DATA/
MOUNTED  HIGH      N          512    4096    4194304  3354624  614328
609930          1466           0          0          N RECO/
MOUNTED  NORMAL   N          512    4096    4194304  894720  893432
81338          406047          0          0          Y DBFS_DG/

```

Note that the DBFS_DG diskgroup contains Voting Files.

- b. As the grid infrastructure owner, verify which diskgroups contain OCR and voting disks

List Voting disks

```

[grid@sclczdb01 ~]$ crsctl query css votedisk
## STATE      File Universal Id                        File Name Disk group
--  -
  1. ONLINE    9e9ed654c2d24f10bfa254f13bc33f49
(o/192.168.75.18/DBFS_DG_CD_02_sclczcel01) [DBFS_DG]
  2. ONLINE    2a38187a65f14f25bf01fe96813fcd80
(o/192.168.75.19/DBFS_DG_CD_02_sclczcel02) [DBFS_DG]
  3. ONLINE    7d0fa058177f4fa7bf0314b4de369257
(o/192.168.75.20/DBFS_DG_CD_02_sclczcel03) [DBFS_DG]

```

Located 3 voting disk(s).

List OCR devices

```
[grid@sclczdb02 ~]$ ocrcheck
Status of Oracle Cluster Registry is as follows :
   Version                    :          3
   Total space (kbytes)       :       262120
   Used space (kbytes)        :         3472
   Available space (kbytes)   :       258648
   ID                          : 1946869935
   Device/File Name           : +DBFS_DG
                               Device/File integrity check succeeded
                               Device/File not configured
                               Device/File not configured
                               Device/File not configured
                               Device/File not configured
```

Cluster registry integrity check succeeded

Logical corruption check bypassed due to non-privileged user

- c. As the grid infrastructure owner, run the following to prepare for dropping any unwanted diskgroups that do not contain OCR or Voting disks.

```
$ srvctl stop diskgroup -g RECO
$ srvctl stop diskgroup -g DATA
$ srvctl remove diskgroup -g DATA -f
$ srvctl remove diskgroup -g RECO -f
```

2. Drop the corresponding griddisks which will also drop the Oracle ASM diskgroups that were stopped and removed from OCR

```
$ dcli -g cell_group -l celladmin cellcli -e drop griddisk all prefix=data
$ dcli -g cell_group -l celladmin cellcli -e drop griddisk all prefix=reco
```

Note: If you receive a "CELL-02550: Cell Server (CELLSRV) cannot drop the grid disk." error then you will have to use the 'FORCE' option. In this case the commands would be:

```
$ dcli -g cell_group -l celladmin cellcli -e drop griddisk all prefix=data
force
$ dcli -g cell_group -l celladmin cellcli -e drop griddisk all prefix=reco
force
```

3. View total available space

```
$ dcli -g cell_group -l celladmin cellcli -e list celldisk where
disktype=HardDisk attributes freeSpace
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel01: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
```

```
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel02: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
sclczcel03: 528.6875G
```

4. Create the data griddisks for each application

Note that the order of creation here is important and that you may even want to specify the `OFFSET` option in the `create griddisk` command. The order is important because the first created griddisks will be on the outer portion of the physical disk and therefore have better performance. If no offset is specified then the next available slice is used and it moves inward on the physical disk.

```
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=DATA_EBS, size=56g"
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=DATA_PSFT, size=91g"
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=DATA_SIEBEL, size=34g"
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=PLACE HOLDER, size=125g"
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=FRA_EBS, size=90g"
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=FRA_PSFT, size=72g"
$ dcli -g cell_group -l celladmin "cellcli -e create griddisk all harddisk
prefix=FRA_SIEBEL, size=60g"
```

Create ASM Diskgroups

These commands should be run as a user with the `asmadmin` role on a single ASM instance. Note that you should check the *Automatic Storage Management Administrator's Guide* ([Table 4-3](#) in the 11.2 guide) for which ASM features are enabled depending on the `COMPATIBLE .ASM` attribute setting.

DATA

a. EBS

```
create diskgroup DATA_EBS high redundancy
disk 'o/*/DATA_EBS*'
```

```

attribute
'cell.smart_scan_capable'='true', 'compatible.rdbms'='11.2', 'compatible.asm
='11.2.0.2.0', 'au_size'='4M';

```

b. PeopleSoft

```

create diskgroup DATA_PSFT high redundancy
disk 'o/*/DATA_PSFT*'
attribute
'cell.smart_scan_capable'='true', 'compatible.rdbms'='11.2', 'compatible.asm
='11.2.0.2.0', 'au_size'='4M';

```

c. Siebel

```

create diskgroup DATA_SIEBEL high redundancy
disk 'o/*/DATA_SIEBEL*'
attribute
'cell.smart_scan_capable'='true', 'compatible.rdbms'='11.2', 'compatible.asm
='11.2.0.2.0', 'au_size'='4M';

```

FRA

a. EBS

```

create diskgroup FRA_EBS normal redundancy
disk 'o/*/FRA_EBS*'
attribute
'cell.smart_scan_capable'='true', 'compatible.rdbms'='11.2', 'compatible.asm
='11.2.0.2.0', 'au_size'='4M';

```

d. Peoplesoft

```

create diskgroup FRA_PSFT normal redundancy
disk 'o/*/FRA_PSFT*'
attribute
'cell.smart_scan_capable'='true', 'compatible.rdbms'='11.2', 'compatible.asm
='11.2.0.2.0', 'au_size'='4M';

```

e. Siebel

```

create diskgroup FRA_SIEBEL normal redundancy
disk 'o/*/FRA_SIEBEL*'
attribute
'cell.smart_scan_capable'='true', 'compatible.rdbms'='11.2', 'compatible.asm
='11.2.0.2.0', 'au_size'='4M';

```

Other Attributes

The ASM disk repair timer represents the amount of time a disk (or failure group, i.e. cell) will remain offline before ASM drops it and runs the subsequent rebalance. While the disk or cell is offline, the redundancy level is not maintained and ASM is tracking the changed extents so the disk can be resynchronized if the problem was temporary. The default disk repair timer is 3.6 hours and can be changed but it should be understood that the redundancy level is not maintained while a disk is offline for the `disk_repair_time` time up until a rebalance operation completes after the `disk_repair_time` has expired. If the default is inadequate, set the disk repair timer to the maximum amount of time it takes to detect and correct a temporary disk failure. This example changes the disk repair timer to 8.5 hours for the diskgroups.

It is important to understand that changing the `disk_repair_timer` parameter does NOT change the repair timer in use for disks currently offline. The repair timer for those offline disks is either

the repair timer specified on the command line (if disks manually offlined) or the default repair timer in effect when the offline occurred. To change the repair timer for disks currently offlined, you can re-issue the offline command with the desired repair timer specified.

It is also important to understand that, typically, disks and cells have different failure and repair characteristics. Typically, a disk failure is permanent (ex: disk dies), and a cell failure is transient (ex: cell reboot). Starting with release 11.2.1.3, a feature was added to proactively drop disks that are predictive failure or failed. This allows the setting of `disk_repair_time` to be geared almost exclusively toward cell failure, thus easing the current aforementioned trade-off. In other words, you can set `disk_repair_time` primarily for cell failure.

```
alter diskgroup DATA_EBS set attribute 'disk_repair_time'='8.5h';
alter diskgroup FRA_EBS set attribute 'disk_repair_time'='8.5h';

alter diskgroup DATA_PSFT set attribute 'disk_repair_time'='8.5h';
alter diskgroup FRA_PSFT set attribute 'disk_repair_time'='8.5h';

alter diskgroup DATA_SIEBEL set attribute 'disk_repair_time'='8.5h';
alter diskgroup FRA_SIEBEL set attribute 'disk_repair_time'='8.5h';
```

Configuring Exadata database-scoped Security Example

Refer to the “Oracle® Exadata Storage Server Software User's Guide, 11g Release 2 (11.2)” [5] chapter 4, “Configuring Security for Oracle Exadata Storage Server Software.” The high-level steps for configuring Exadata database-scoped security are:

1. [Setup ASM-scoped Security](#)
2. [Configure database-scoped Security](#)
3. [Validate the Configuration](#)

Setup ASM-scoped Security

1. Get the ASM database `DB_UNIQUE_NAME`

```
[grid@sclczdb01 ~]$ sqlplus '/ as sysasm'
SQL> show parameter uniq
NAME                                TYPE        VALUE
-----
db_unique_name                       string      +ASM
```

Note: The `db_unique_name` is case-sensitive. Ensure the case matches for all subsequent settings

2. Shut down the database and Oracle ASM instances that will have their security configuration changed.

```
# /u01/app/11.2.0/grid/bin/crsctl stop crs
```

3. Login on a cell as celladmin, run cellcli and use the cellcli CREATE KEY command to generate a random hexadecimal string. The command can be run on any cell. After running the command, the system will display the new key.

```
CellCLI> create key
5735d31101dfcefcd8ca077afddf7c
```

4. Set the cellkey.ora file up as the Oracle Grid Infrastructure owner

```
[grid@sclczdb01 ~]$ cat /home/grid/cellkey.ora
key=5735d31101dfcefcd8ca077afddf7c
asm=+ASM
#realm=my_realm
```

5. Use the ASSIGN KEY command to assign the security key to the Oracle ASM cluster client **on all the cells** that you want the Oracle ASM cluster to access.

Use the DB_UNIQUE_NAME captured earlier. Note that this is case-sensitive.

```
# dcli -g ~/cell_group -l celladmin cellcli -e ASSIGN KEY FOR
+ASM='5735d31101dfcefcd8ca077afddf7c'
sclczcel01: Key for +ASM successfully created
sclczcel02: Key for +ASM successfully created
sclczcel03: Key for +ASM successfully created
```

Note: if you rerun this then it alters the cell to the new key

6. Enter the Oracle ASM DB_UNIQUE_NAME in the availableTo attribute of each griddisk with the CREATE GRIDDISK or ALTER GRIDDISK command to configure security on the grid disks on all the cells that you want the Oracle ASM cluster to access.

```
# dcli -g ~/cell_group -l celladmin "cellcli -e ALTER GRIDDISK ALL
availableTo='\'+ASM\'"
```

7. Validate the change

```
# dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes
availableto" | sort -u
sclczcel01: +ASM
sclczcel02: +ASM
sclczcel03: +ASM
```

Setup the ASM Scoped-security Clients

1. Copy the cellkey.ora file created above to each compute node's /etc/oracle/cell/network-config directory. The /etc/oracle/cell/network-config/cellkey.ora file is required for ASM-scoped security only.

```
[grid@sclczdb01 ~]$ dcli -g ~/dbs_group -l grid -f /home/grid/cellkey.ora -d
/etc/oracle/cell/network-config
```

2. Set the file privileges

```
[grid@sclczdb01 ~]$ dcli -g ~/dbs_group -l grid chmod 600
/etc/oracle/cell/network-config/cellkey.ora
```

```
[grid@sclczdb01 ~]$ dcli -g ~/dbs_group -l grid ls -l
/etc/oracle/cell/network-config/cellkey.ora
sclczdb01: -rw----- 1 grid oinstall 62 Mar 11 13:38
/etc/oracle/cell/network-config/cellkey.ora
```

```
sclczdb02: -rw----- 1 grid oinstall 62 Mar 11 13:38
/etc/oracle/cell/network-config/cellkey.ora
```

- Restart the instances after you have created the cellkey.ora files.

```
# /u01/app/11.2.0/grid/bin/crsctl start crs
```

- Verify that all resources are online

```
[grid@sclczdb01 ~]$ /u01/app/11.2.0/grid/bin/crs_stat -t
```

- Remove the temporary cell.key file

```
[grid@sclczdb01 ~]$ rm /home/grid/cellkey.ora
```

Configure database-scoped Security

Note: The db_unique_name is case-sensitive. Ensure the case matches for all subsequent settings

Note that database-scoped security can be setup under a single operating system user but since our case study requires to isolate the database administration for each database we are setting up separate operating system accounts for each application database and each will have its own database-scoped security key and own set of diskgroups. This example is dependent on each application operating system account having ssh equivalence setup to the compute nodes. If that is not setup then you would have to run all commands from the root user and then change the cellkey.ora file owner to the application owner and group. For example:

```
# chown oracle_ebs:oinstall
/u01/app/oracle_ebs/product/11.2.0.2/VIS_RAC/admin/VIS/pfile/cellkey.ora
```

- Shut down the database and Oracle ASM instances that will have their security configuration changed.

```
/u01/app/11.2.0/grid/bin/crsctl stop crs
```

- Login on a cell as celladmin, run cellcli and use the cellcli CREATE KEY command to generate a random hexadecimal string for each application database that you want to access specific grid disks. The command can be run on any cell. After running the command, the system will display the new key. These keys are used to configure the database client access and also assigned to each Exadata cell with a database name as illustrated in subsequent steps.

```
CellCLI> create key
7308c1ad5e2fed442b849c1487f94c4
CellCLI> create key
17ab0274cac3b10d2fd42dfa2744a9c0
CellCLI> create key
1e7237d8a8f679753ef6c9d5acbd477b
These keys are used to configure
```

- Set the cellkey.ora file up for each application database owner
 - Oracle E-Business Suite (db_unique_name=VIS)

```
[oracle_ebs@sclczdb01 ~]$ cat cellkey.ora
key=7308c1ad5e2fed442b849c1487f94c4
asm+=ASM
```



```
[oracle_ebs@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_ebs mkdir -p
$ORACLE_HOME/admin/VIS/pfile

[oracle_ebs@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_ebs -f cellkey.ora -
d $ORACLE_HOME/admin/VIS/pfile

[oracle_ebs@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_ebs chmod 600
$ORACLE_HOME/admin/VIS/pfile/cellkey.ora

[oracle_ebs@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_ebs ls -l
$ORACLE_HOME/admin/VIS/pfile/cellkey.ora
sclczdb01: -rw----- 1 oracle_ebs oinstall 62 Mar 15 21:42
/u01/app/oracle_ebs/product/11.2.0.2/VIS_RAC/admin/VIS/pfile/cellkey.ora
sclczdb02: -rw----- 1 oracle_ebs oinstall 62 Mar 15 21:42
/u01/app/oracle_ebs/product/11.2.0.2/VIS_RAC/admin/VIS/pfile/cellkey.ora
```

b. Oracle Peoplesoft (db_unique_name =psft)

```
[oracle_psft@sclczdb01 ~]$ cat cellkey.ora
key=17ab0274cac3b10d2fd42dfa2744a9c0
asm=+ASM
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_psft mkdir -p
$ORACLE_HOME/admin/psft/pfile
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_psft -f
cellkey.ora -d $ORACLE_HOME/admin/psft/pfile
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_psft chmod 600
$ORACLE_HOME/admin/psft/pfile/cellkey.ora
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_psft ls -l
$ORACLE_HOME/admin/psft/pfile
sclczdb01: total 4
sclczdb01: -rw----- 1 oracle_psft oinstall 46 Mar 15 21:18 cellkey.ora
sclczdb02: total 4
sclczdb02: -rw----- 1 oracle_psft oinstall 46 Mar 15 21:18 cellkey.ora
```

c. Siebel (db_unique_name =quarter)

```
[oracle_siebel@sclczdb01 ~]$ cat cellkey.ora
key=1e7237d8a8f679753ef6c9d5acbd477b
asm=+ASM
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_siebel mkdir -p
$ORACLE_HOME/admin/quarter/pfile
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_siebel -f
cellkey.ora -d $ORACLE_HOME/admin/quarter/pfile
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_siebel chmod 600
$ORACLE_HOME/admin/quarter/pfile/cellkey.ora
[oracle_siebel@sclczdb01 ~]$ dcli -g ~/dbs_group -l oracle_siebel ls -l
$ORACLE_HOME/admin/quarter/pfile
sclczdb01: total 4
sclczdb01: -rw-r----- 1 oracle_siebel oinstall 46 Mar 14 18:10 cellkey.ora
sclczdb02: total 4
sclczdb02: -rw-r----- 1 oracle_siebel oinstall 46 Mar 14 18:10 cellkey.ora
```

4. Use the ASSIGN KEY command to assign the keys to database clients on the cells that contain the grid disks

```
# dcli -g ~/cell_group -l root "cellcli -e ASSIGN KEY FOR
VIS=\`7308clad5e2fed442b849c1487f94c4\`,psft=\`17ab0274cac3b10d2fd42dfa2744a
9c0\`,quarter=\`1e7237d8a8f679753ef6c9d5acbd477b\`"
sclczcel01: Key for VIS successfully created
sclczcel01: Key for psft successfully created
sclczcel01: Key for quarter successfully created
sclczcel02: Key for VIS successfully created
sclczcel02: Key for psft successfully created
sclczcel02: Key for quarter successfully created
```

```
sclczcel03: Key for VIS successfully created
sclczcel03: Key for psft successfully created
sclczcel03: Key for quarter successfully created
```

```
# dcli -g ~/cell_group -l celladmin cellcli -e LIST KEY
sclczcel01: +ASM          5735d31101dfcefcd8ca077afddd7c
sclczcel01: psft         17ab0274cac3b10d2fd42dfa2744a9c0
sclczcel01: quarter     1e7237d8a8f679753ef6c9d5acbd477b
sclczcel01: VIS         7308clad5e2fed442b849c1487f94c4
sclczcel02: +ASM          5735d31101dfcefcd8ca077afddd7c
sclczcel02: psft         17ab0274cac3b10d2fd42dfa2744a9c0
sclczcel02: quarter     1e7237d8a8f679753ef6c9d5acbd477b
sclczcel02: VIS         7308clad5e2fed442b849c1487f94c4
sclczcel03: +ASM          5735d31101dfcefcd8ca077afddd7c
sclczcel03: psft         17ab0274cac3b10d2fd42dfa2744a9c0
sclczcel03: quarter     1e7237d8a8f679753ef6c9d5acbd477b
sclczcel03: VIS         7308clad5e2fed442b849c1487f94c4
```

5. Set the availableTo attribute on each griddisk with the CREATE GRIDDISK or ALTER GRIDDISK command to configure security on grid disks. You must include the Oracle ASM cluster name with the database clients names when setting the value of the availableTo attribute. The names that you enter must be the Oracle ASM database and the Oracle RAC cluster unique names (DB_UNIQUE_NAME).

This command generates a script to alter each griddisk for the EBS diskgroups. This can be executed for each application that has its own diskgroups making sure to change the **griddisk WHERE clause** and the DB_UNIQUE_NAME.

- a. Generate a script to alter each griddisk's availableTo attribute for the EBS diskgroups.

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk where name
like \'.*_EBS_*\' | awk -F: '{print $1 " " $2}' | awk '{print "dcli -
c " $1 " -l celladmin \"cellcli -e alter griddisk " $2 " availableTo="
"\x5c\x27" "+ASM,VIS" "\x5c\x27" "\"}' > dbScoped_EBS.sh
```

```
# chmod 755 dbScoped_EBS.sh
```

```
# ./dbScoped_EBS.sh
```

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes
name,availableTo where name like \'.*_EBS_*\'"
```

- b. Generate a script to alter each griddisk's availableTo attribute for the Peoplesoft diskgroups.

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk where name
like \'.*_PSFT_*\' | awk -F: '{print $1 " " $2}' | awk '{print "dcli -
c " $1 " -l celladmin \"cellcli -e alter griddisk " $2 " availableTo="
"\x5c\x27" "+ASM,psft" "\x5c\x27" "\"}' > dbScoped_PSFT.sh
```

```
# chmod 755 dbScoped_PSFT.sh
```

```
# ./dbScoped_PSFT.sh
```

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes
name,availableTo where name like \'.*_PSFT_.*\'"
```

- c. Generate a script to alter each griddisk's availableTo attribute for the Siebel diskgroups.

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk where name
like \'.*_SIEBEL_.*\'" | awk -F: '{print $1 " " $2}' | awk '{print "dcli
-c " $1 " -l celladmin \"cellcli -e alter griddisk " $2 " availableTo="
"\x5c\x27" "+ASM,quarter" "\x5c\x27" "\""}' > dbScoped_SIEBEL.sh
```

```
# chmod 755 dbScoped_SIEBEL.sh
```

```
# ./dbScoped_SIEBEL.sh
```

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes
name,availableTo where name like \'.*_SIEBEL_.*\'"
```

- d. Validate the availableTo attribute setting is correct

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes
availableTo" | sort -u
```

```
sclczcel01: +ASM
sclczcel01: +ASM,psft
sclczcel01: +ASM,quarter
sclczcel01: +ASM,VIS
sclczcel02: +ASM
sclczcel02: +ASM,psft
sclczcel02: +ASM,quarter
sclczcel02: +ASM,VIS
sclczcel03: +ASM
sclczcel03: +ASM,psft
sclczcel03: +ASM,quarter
sclczcel03: +ASM,VIS
```

- e. Restart the instances after you have created and edited the cellkey.ora files.

```
# /u01/app/11.2.0/grid/bin/crsctl start crs
```

Validate the Configuration

1. Show ASM level key files

```
# dcli -g ~/dbs_group -l root ls -l /etc/oracle/cell/network-
config/cellkey.ora
```

```
sclczdb01: -rw----- 1 grid oinstall 62 Mar 14 15:24
/etc/oracle/cell/network-config/cellkey.ora
sclczdb02: -rw----- 1 grid oinstall 62 Mar 14 15:24
/etc/oracle/cell/network-config/cellkey.ora
```

```
# dcli -g ~/dbs_group -l root cat /etc/oracle/cell/network-config/cellkey.ora
```

```
sclczdb01: key=5735d31101dfcefcdb8ca077afddd7c
sclczdb01: asm=+ASM
sclczdb01: #realm=my_realm
sclczdb02: key=5735d31101dfcefcdb8ca077afddd7c
sclczdb02: asm=+ASM
```

```
sclczdb02: #realm=my_realm
```

2. Show DB level key files

```
# dcli -g ~/dbs_group -l root ls -l
/u01/app/*/product/11.2.0.2/admin/*/pfile/*
sclczdb01: -rw----- 1 oracle_ebs      oinstall 62 Mar 15 21:42
/u01/app/oracle_ebs/product/11.2.0.2/VIS_RAC/admin/VIS/pfile/cellkey.ora
sclczdb01: -rw----- 1 oracle_psft     oinstall 46 Mar 15 21:18
/u01/app/oracle_psft/product/11.2.0.2/dbhome_psft/admin/PSFT/pfile/cellkey.ora
sclczdb01: -rw----- 1 oracle_siebel  oinstall 46 Mar 16 11:33
/u01/app/oracle_siebel/product/11.2.0/dbhome_siebel/admin/quarter/pfile/cellkey.ora
sclczdb02: -rw----- 1 oracle_ebs      oinstall 62 Mar 15 21:42
/u01/app/oracle_ebs/product/11.2.0.2/VIS_RAC/admin/VIS/pfile/cellkey.ora
sclczdb02: -rw----- 1 oracle_psft     oinstall 46 Mar 15 21:18
/u01/app/oracle_psft/product/11.2.0.2/dbhome_psft/admin/PSFT/pfile/cellkey.ora
sclczdb02: -rw----- 1 oracle_siebel  oinstall 46 Mar 16 11:33
/u01/app/oracle_siebel/product/11.2.0/dbhome_siebel/admin/quarter/pfile/cellkey.ora
```

```
# dcli -g ~/dbs_group -l root cat
/u01/app/*/product/11.2.0.2/admin/*/pfile/cellkey.ora
sclczdb01: key=7308c1ad5e2fed442b849c1487f94c4
sclczdb01: asm=+ASM
sclczdb01: #realm=my_realm
sclczdb01:
sclczdb01: key=17ab0274cac3b10d2fd42dfa2744a9c0
sclczdb01: asm=+ASM
sclczdb01: key=1e7237d8a8f679753ef6c9d5acbd477b
sclczdb01: asm=+ASM
sclczdb02: key=7308c1ad5e2fed442b849c1487f94c4
sclczdb02: asm=+ASM
sclczdb02: #realm=my_realm
sclczdb02:
sclczdb02: key=17ab0274cac3b10d2fd42dfa2744a9c0
sclczdb02: asm=+ASM
sclczdb02: key=1e7237d8a8f679753ef6c9d5acbd477b
sclczdb02: asm=+ASM
```

3. Show griddisk attributes

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes \
name,availableTo where name like \'.*_EBS_.*\' | sort -u
sclczcel01: +ASM,VIS
sclczcel02: +ASM,VIS
sclczcel03: +ASM,VIS
```

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes \
availableTo where name like \'.*_PSFT_.*\' | sort -u
sclczcel01: +ASM,PSFT
sclczcel02: +ASM,PSFT
sclczcel03: +ASM,PSFT
```

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list griddisk attributes \
availableTo where name like \'.*_SIEBEL_.*\' | sort -u
sclczcel01: +ASM,quarter
sclczcel02: +ASM,quarter
sclczcel03: +ASM,quarter
```

4. Show cell key assignments

```
$ dcli -g ~/cell_group -l celladmin "cellcli -e list key"
sclczcel01: +ASM          5735d31101dfcefcd8ca077afddf7c
sclczcel01: PSFT         17ab0274cac3b10d2fd42dfa2744a9c0
sclczcel01: quarter      1e7237d8a8f679753ef6c9d5acbd477b
```

```

sclczcel01: VIS          7308c1ad5e2fed442b849c1487f94c4
sclczcel02: +ASM        5735d31101dfcefcdb8ca077afdddf7c
sclczcel02: PSFT        17ab0274cac3b10d2fd42dfa2744a9c0
sclczcel02: quarter     1e7237d8a8f679753ef6c9d5acbd477b
sclczcel02: VIS          7308c1ad5e2fed442b849c1487f94c4
sclczcel03: +ASM        5735d31101dfcefcdb8ca077afdddf7c
sclczcel03: PSFT        17ab0274cac3b10d2fd42dfa2744a9c0
sclczcel03: quarter     1e7237d8a8f679753ef6c9d5acbd477b
sclczcel03: VIS          7308c1ad5e2fed442b849c1487f94c4

```

dbScoped_EBS.sh Generated Script

This script is generated to set the availableTo attribute for each griddisk for an application's griddisks as part of the [database-scoped Security Setup](#).

```

dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_00_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_01_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_02_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_03_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_04_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_05_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_06_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_07_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_08_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_09_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_10_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_11_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_00_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_01_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_02_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_03_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_04_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_05_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_06_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_07_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_08_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_09_sclczcel01 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_10_sclczcel01 availableTo=\'+ASM,VIS\'"

```

```

dcli -c sclczcel01 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_11_sclczcel01 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_00_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_01_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_02_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_03_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_04_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_05_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_06_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_07_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_08_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_09_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_10_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_11_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_00_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_01_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_02_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_03_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_04_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_05_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_06_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_07_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_08_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_09_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_10_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel02 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_11_sclczcel02 availableTo='+ASM,VIS'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_00_sclczcel03 availableTo='+ASM,VIS'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_01_sclczcel03 availableTo='+ASM,VIS'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_02_sclczcel03 availableTo='+ASM,VIS'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_03_sclczcel03 availableTo='+ASM,VIS'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_04_sclczcel03 availableTo='+ASM,VIS'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_05_sclczcel03 availableTo='+ASM,VIS'"

```

```
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_06_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_07_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_08_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_09_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_10_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
DATA_EBS_CD_11_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_00_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_01_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_02_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_03_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_04_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_05_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_06_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_07_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_08_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_09_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_10_sclczcel03 availableTo=\'+ASM,VIS\'"
dcli -c sclczcel03 -l celladmin "cellcli -e alter griddisk
FRA_EBS_CD_11_sclczcel03 availableTo=\'+ASM,VIS\'"
```

References

1. Oracle Maximum Availability Architecture Web site
<http://www.otn.oracle.com/goto/maa>
2. [Oracle Database Security Checklist](#) - Technical Whitepaper
3. [Oracle® Grid Infrastructure Installation Guide 11g Release 2 \(11.2\) for Linux - Creating Groups, Users and Paths for Oracle Grid Infrastructure](#)
4. [Oracle Automatic Storage Management Administrator's Guide 11g Release 2 \(11.2\)](#) – Chapter 3, [Using Separate Operating System Groups for Oracle ASM Users](#)
5. Oracle Exadata Storage Server Software User's Guide, 11g Release 2 (11.2)
See /opt/oracle/cell/doc/doc on any Exadata Storage Server
6. Oracle White Paper - [Cost Effective Security and Compliance with Oracle Database 11g Release 2](#)
7. [Oracle® Database Security Guide, 11g Release 2 \(11.2\)](#)
8. [Oracle® Database 2 Day + Security Guide, 11g Release 2 \(11.2\)](#)
9. [Special Publication 800-12: An Introduction to Computer Security - The NIST Handbook](#)
10. Oracle Exadata Database Machine Owner's Guide, 11g Release 2 (11.2)
See /opt/oracle/cell/doc/doc on any Exadata Storage Server
11. [MAA Best Practices for Oracle Exadata Database Machine \(technical white paper\)](#)

Change Record

Date	Summary of Changes
7/20/12	<ul style="list-style-type: none"> • Revised the diskgroup creation examples to set COMPATIBLE.RDBMS='11.2' instead of 11.2.0.1. See the ‘Create ASM Diskgroups’ section.
4/18/12	<ul style="list-style-type: none"> • Revised table formats to not split rows across a page and repeat headings • Removed mkdir \$username line from the setupOS*Users.sh scripts • Added ‘-f’ to the ‘srvctl remove diskgroup -g RECO -f’ command under the “Griddisk and ASM Diskgroup Creation” / “Create Griddisks” section in step 1c. Also added a note regarding the possibility of needing the FORCE option on the ‘drop griddisk’ command in step 2.
9/12/11	<ul style="list-style-type: none"> • Added example output for some examples • Added “chown -R \$username:oinstall /home/\$username” to the setupOS_dbaUsers.sh script under “Create Post-install OS Accounts”. • Added the asmdba OS group to the useradd command example • Changed diskgroup removal commands to use srvctl under “Griddisk and ASM Diskgroup Creation”

Date	Summary of Changes
8/31/11	Added a note on the ASM disk group limit of 63 per storage system under " Diskgroup Isolation ".
8/24/11	Contributor name correction
8/19/11	Review complete for publishing
8/5/11	Second internal review cycle
7/18/11	Initial internal review



Oracle Exadata Database Machine
Consolidation: Segregating Databases and
Roles
July 2012
Author: Ray Dutcher
Contributing Authors: Richard Exley, Richard
Jobin, Darryl Presley, Lyn Pratt, Dan Norris,
MAA Team

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



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

Copyright © 2011, 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.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together