

BIAS

Oracle Exadata X6 Pallet to Production:

An All-Inclusive Migration Package

BIAS



BIAS
STRENGTH IN DIVERSITY

Company Overview

Our Team

BIAS

About Us

Locations



Founded in 2000



2018 Oracle Excellence Award Winner



~400 employees with 10+ years of Oracle experience on average



Oracle Platinum Partner
Cloud Excellence Implementer



Service Organization Control (SOC) 2 Certified



35 Oracle Specializations spanning the entire stack



Inc. 500 | 5000 Fastest Growing Private Companies in the U.S. for the 8th Time



2018 Best Places to Work

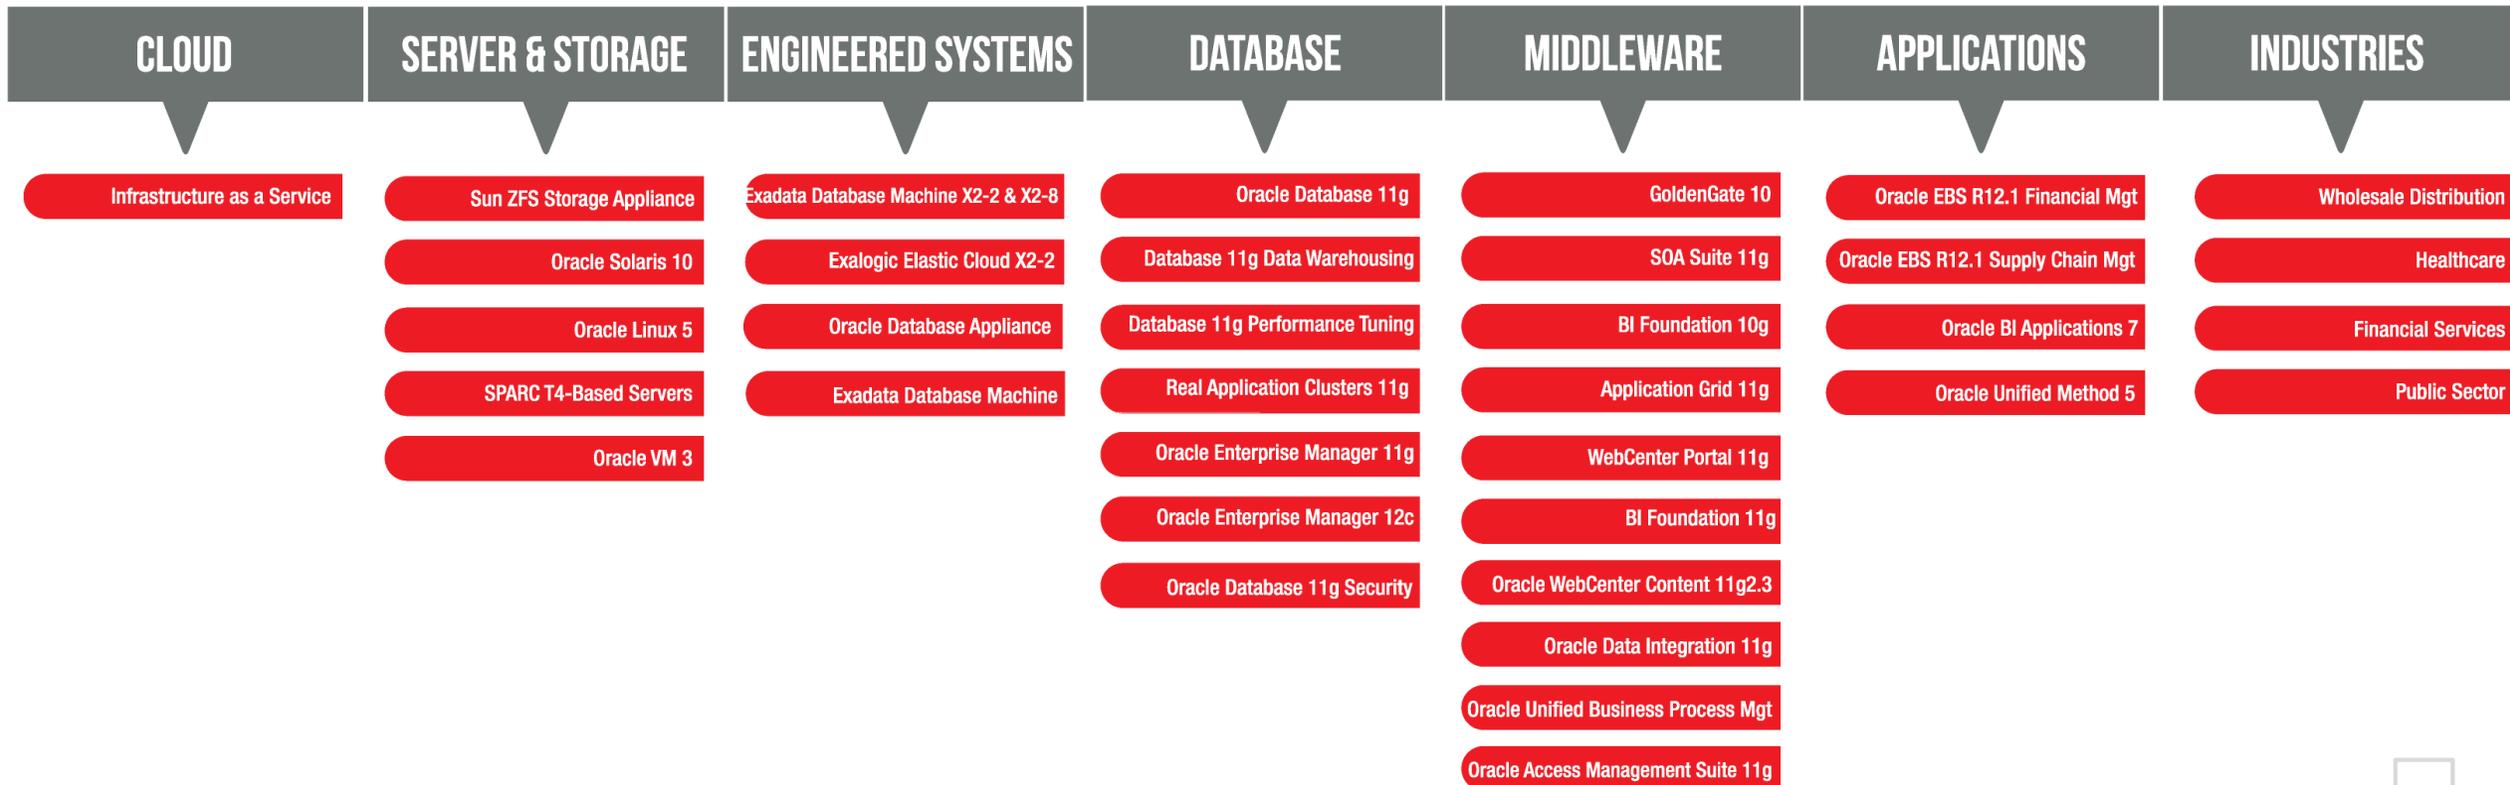


ISO 27001 Certified - Information Security Management System

BIAS Oracle Specializations



Oracle created the OPN Specialized Program to showcase the Oracle partners who have achieved expertise in Oracle product areas and reached specialization status through competency development, business results, expertise and proven success. BIAS is proud to be specialized in 35 areas of Oracle products, which include the following:



BIAS



AGENDA

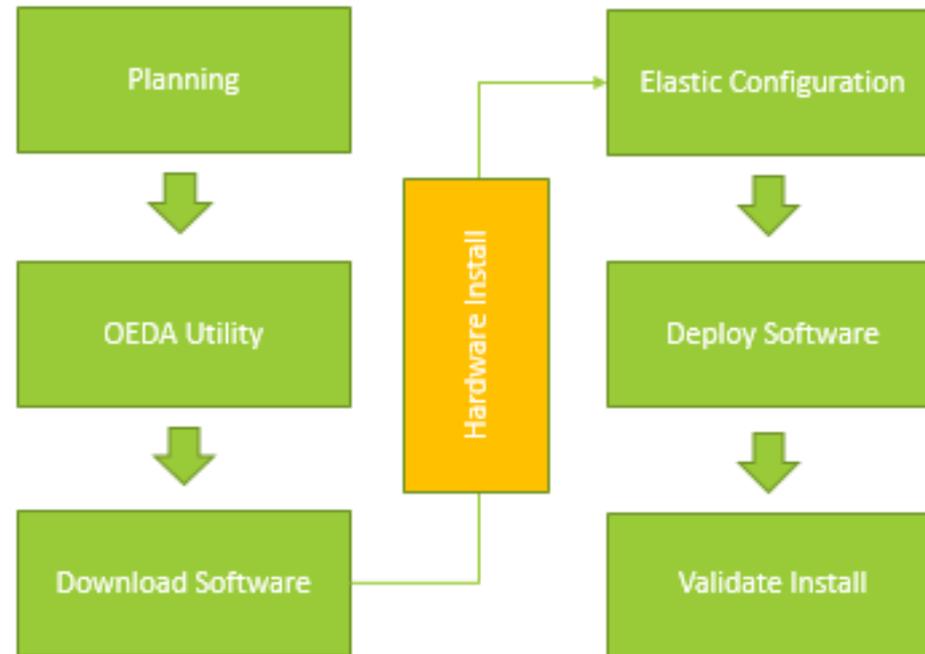
BIAS

- 1 **Exadata Configuration**
- 2 **Disk Group Sizing**
- 3 **Current & Future State Architecture**
- 4 **Migrations**
- 5 **Migration Issues**
- 6 **Exadata Snap Clone**

Exadata Configuration Overview



Overview



Exadata Configuration Considerations

BIAS

Operating System

- Networking
 - DNS
 - NTP
 - Gateway
 - Client
 - Management
 - Infiniband
 - Backup
- Host Naming Conventions
- Capacity on Demand (Licensing)
- Virtualization
- Role Separation
- Group and User IDs
- Time Zone

Versions

- Image
- Clusterware
- RDBMS

+ASM Storage

- Naming
- Redundancy
- Layout
- Sparse Disk

Exadata Configuration OEDA



☆ Exadata Database Machine and Exadata Storage Server Supported Versions (Doc ID 88828.1)

In this Document

[Purpose](#)

[Scope](#)

[Details](#)

[Latest Releases and Patching News](#)

[Exadata Software Updates Overview and Guidelines](#)

[Exadata Software and Hardware Maintenance Planning](#)

[Critical Issues](#)

[Security-Related Guidance](#)

Latest Releases

[Patch Release Frequency](#)

[Recommended Versions](#)

[Latest Releases](#)

[Guidelines for Related Products and Features](#)

[Version Compatibility and Requirements](#)

[Exadata Database Machine Hardware](#)

[Oracle Database](#)

[Minimum Exadata, Grid Infrastructure, and Database Requirements for Exadata Feature Usage](#)

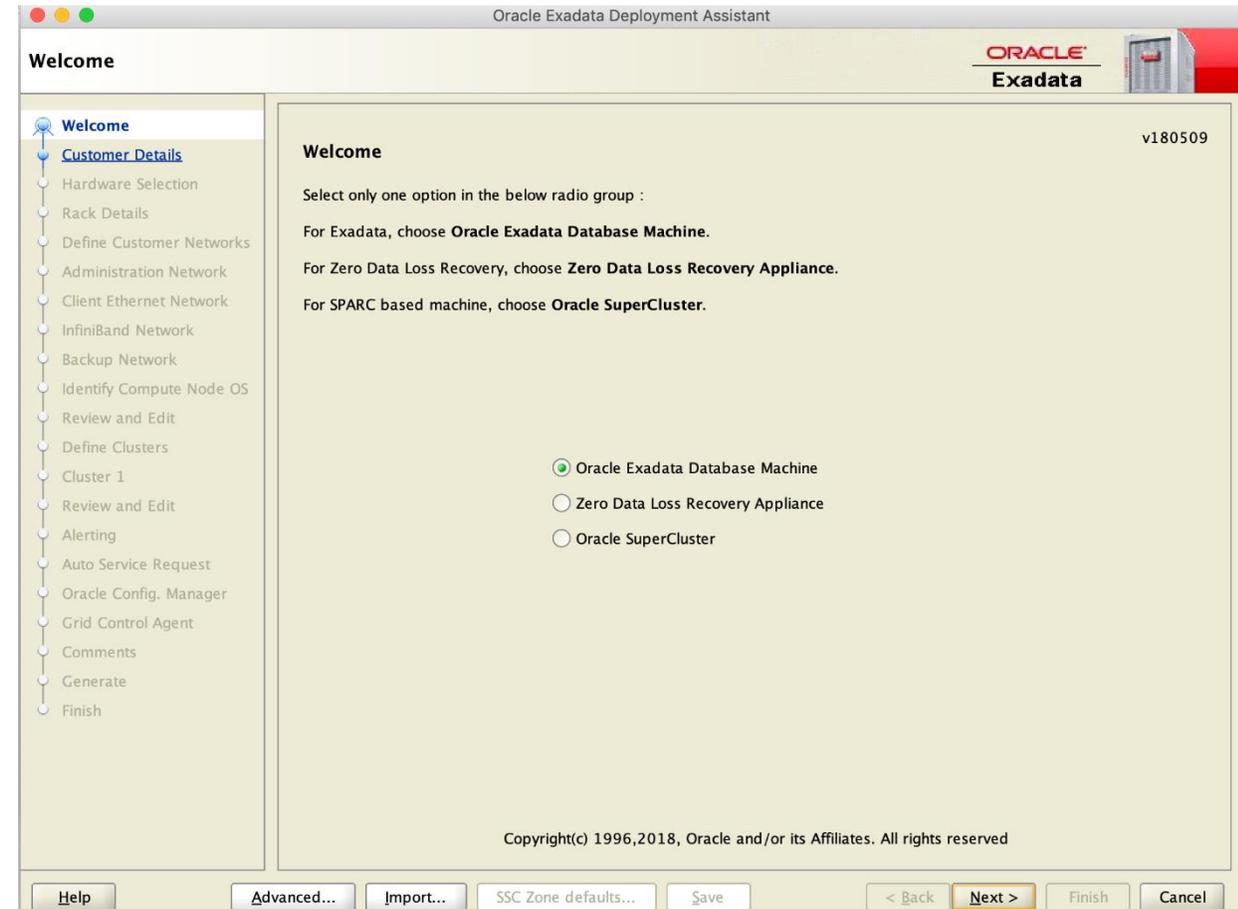
[Database Server Operating System](#)

[Oracle Linux x86-64 on 2-socket Intel x86-64 Database Servers](#)

[Oracle Linux x86-64 on 8-socket Intel x86-64 Database Servers](#)

[Oracle Solaris on 2-socket \(X4-2, X3-2, X2-2\) Database Servers](#)

[Patch Release History for Exadata Database Machine Components](#)



Exadata Configuration Networking

BIAS

Oracle Exadata Deployment Assistant - Customer Details screen. The interface shows a sidebar with navigation options: Welcome, Customer Details, Hardware Selection, Rack Details, Define Customer Networks, Administration Network, Client Ethernet Network, InfiniBand Network, Backup Network, Identify Compute Node OS, Review and Edit, Define Clusters, Cluster 1, Review and Edit, Alerting, Auto Service Request, Oracle Config. Manager, Grid Control Agent, Comments, Generate, and Finish. The main area is titled 'Customer Details' and contains the following fields:

- Customer Name: Genuine Parts Company
- Application: GPC
- Network Domain Name: goname.net
- Name Prefix: Gexa
- Region: America
- Timezone: America/New_York
- DNS: 101.6.152.5
- NTP: 101.6.152.5

Buttons at the bottom include: Help, Advanced..., Import..., Oracle VM Defaults..., Save, < Back, Next >, Finish, and Cancel.

Oracle Exadata Deployment Assistant - Define Customer Networks screen. The sidebar is the same as the previous screen. The main area is titled 'Define Customer Networks' and includes a note: 'Exadata requires a minimum of 2 separate customer subnets. This page allows you to describe those subnets, for completeness it also includes subnet 3 which is the Private infiniband network. Some customers have more than 2 subnets. In those cases Exadata can configure one of those additional subnets for "backup", "replication", "dir" or for an "independent client" network in multi cluster environments. This is included here as subnet 4 however this subnet is NOT mandatory for deployment - Click Advanced button to enable infiniband security and VLAN setting'. Below this are four subnets:

- Subnet 1:** Name: Admin, Subnet Mask: 255.255.255.0, Gateway: 101.6.158.1, Admin Network Format: 1/10 Gbit Copper Base-T.
- Subnet 2:** Name: Client, Subnet Mask: 255.255.255.0, Gateway: 101.6.132.1, Client Network Format: 10 Gbit optical.
- Subnet 3:** Name: Private, Subnet Mask: 255.255.255.0, Private Network Format: InfiniBand.
- Subnet 4:** Available Network: Backup, Subnet Mask: 255.255.255.0, Gateway: 101.6.20.1, Backup Network Format: 1/10 Gbit Copper Base-T.

Buttons at the bottom include: Help, Advanced..., Import..., Oracle VM Defaults..., Save, < Back, Next >, Finish, and Cancel.

Oracle Exadata Deployment Assistant - Administration Network screen. The sidebar is the same as the previous screens. The main area is titled 'Administration Network' and contains the following fields:

- Starting IP Address for Pool: 101.6.158.37 (Valid network range: 101.6.158.2 - 101.6.158.254)
- Pool Size: 15
- Ending IP Address for Pool: 101.6.158.51
- Is the default gateway for database servers:
- Defines the Hostname for the database servers:
- Database Server Admin Name: dexadb01, ILOM Name: dexadb01-ilm
- Storage Server Admin Name: dexaceladm01, ILOM Name: dexaceladm01-ilm

Buttons at the bottom include: Modify...

Oracle Exadata Deployment Assistant - InfiniBand Network screen. The sidebar is the same as the previous screens. The main area is titled 'InfiniBand Network' and contains the following fields:

- Starting IP Address for Pool: 192.168.10.10
- Pool Size: 10
- Ending IP Address for Pool: 192.168.10.19
- Enable Active Bonding on Compute node Network:
- Compute Priv Name: dexadb01-priv, Cell Priv Name: dexacel01-priv

Buttons at the bottom include: Modify...

Oracle Exadata Deployment Assistant - Client Ethernet Network screen. The sidebar is the same as the previous screens. The main area is titled 'Client Ethernet Network' and contains the following fields:

- Starting IP Address for Pool: 101.6.132.37 (Valid network range: 101.6.132.2 - 101.6.132.254)
- Pool Size: 7
- Ending IP Address for Pool: 101.6.132.43
- Is the default gateway for database servers:
- Defines the hostname for the database servers:
- Compute Client Name: dexadb01, VIP Name: dexadb01-vip
- Client Scan Name: dexa-scan

Buttons at the bottom include: Modify...

Exadata Configuration Networking & OS

BIAS

Oracle Exadata Deployment Assistant

Review and Edit

ORACLE Exadata

Welcome

Customer Details

Hardware Selection

Rack Details

Define Customer Networks

Administration Network

Client Ethernet Network

InfiniBand Network

Backup Network

Identify Compute Node OS

Review and Edit

Define Clusters

Cluster 1

Review and Edit

Alerting

Auto Service Request

Oracle Config. Manager

Grid Control Agent

Comments

Generate

Finish

Management and Private Networks

This page captures node specific data for the Management, ILOM and Private Networks for the Compute Nodes, Storage Cells and the switches used in the Rack. The Client, VIP, SCAN and backup network names/IP address are collected later in the interview process. For the Capacity-on-Demand feature you can specify the number of enabled cores per physical node

* Indicates host name or IP address that could not be resolved, and the Lookup IP button may not provide immediate feedback to the screen reader

Re-Generate Data Lookup IP

Rack 1

Exadata X6-2 Compute Node 1 Enabled Cores: 36 Rack 1 - Rack Location 16

Admin Name: dexadb01.noname.net
Admin IP: 101.6.158.37
IloM Name: dexadb01-ilon.noname.net
IloM IP: 101.6.158.42
Priv Name: dexadb01-priv1.noname.net Priv IP: 192.168.10.10

Exadata X6-2 Compute Node 2 Enabled Cores: 36 Rack 1 - Rack Location 17

Admin Name: dexadb02.noname.net
Admin IP: 101.6.158.38
IloM Name: dexadb02-ilon.noname.net
IloM IP: 101.6.158.43
Priv Name: dexadb02-priv1.noname.net Priv IP: 192.168.10.12

Exadata Cell Node HC 8TB 1 Rack 1 - Rack Location 2

Admin Name: dexaceladm01.noname.net
Admin IP: 101.6.158.39
IloM Name: dexaceladm01-ilon.noname.net
IloM IP: 101.6.158.44
Priv Name: dexacel01-priv1.noname.net Priv IP: 192.168.10.14

Exadata Cell Node HC 8TB 2 Rack 1 - Rack Location 4

Admin Name: dexaceladm02.noname.net
Admin IP: 101.6.158.40
IloM Name: dexaceladm02-ilon.noname.net
IloM IP: 101.6.158.45
Priv Name: dexacel02-priv1.noname.net Priv IP: 192.168.10.16

Exadata Cell Node HC 8TB 3 Rack 1 - Rack Location 6

Admin Name: dexaceladm03.noname.net
Admin IP: 101.6.158.41
IloM Name: dexaceladm03-ilon.noname.net
IloM IP: 101.6.158.46
Priv Name: dexacel03-priv1.noname.net Priv IP: 192.168.10.18

Exadata Admin Switch 1 Rack 1 - Rack Location 21

Admin Name: dexasw-adm01.noname.net
Admin IP: 101.6.158.47

Exadata Leaf Switch 1 Rack 1 - Rack Location 20

Admin Name: dexasw-lba01.noname.net
Admin IP: 101.6.158.48

Exadata Leaf Switch 2 Rack 1 - Rack Location 22

Admin Name: dexasw-lbb01.noname.net
Admin IP: 101.6.158.49

Exadata PDU Switch 1 Rack 1 - Rack Location 0

Admin Name: dexasw-pdua01.noname.net
Admin IP: 101.6.158.50

Exadata PDU Switch 2 Rack 1 - Rack Location 0

Admin Name: dexasw-pdub01.noname.net
Admin IP: 101.6.158.51

Help Advanced... Import... Oracle YW Defaults... Save < Back Next > Finish Cancel

Users and Groups

Role Separated

Grid ASM Home OS User: grid ID: 100 base: /u01/app/grid

ASM DBA Group: asmdba ID: 1050

ASM Home Oper Group: asmoper ID: 1060

ASM Home Admin Group: asmadmin ID: 1070

RDBMS Home OS User: oracle ID: 1030 base: /u01/app/oracle

RDBMS DBA Group: dba ID: 1005

RDBMS Home Oper Group: racoper ID: 1040

OINSTALL Group: oinstall ID: 1010

Software Locations

Inventory Location: /u01/app/orainventory

Grid Infrastructure Home: 18.2.0.0 RU180417 /u01/app/18.0.0.0/grid

Database Home Location: 12.1.0.2 DBBP170418 /u01/app/oracle/product/12.1.0.2/dbhome_1

Exadata Configuration Networking Summary

Management Network Configuration

All IP addresses and corresponding host names should be registered in DNS
The Admin network defines the hostname and default gateway for all Storage nodes
The Admin network defines the hostname for all Compute nodes

Starting IP Address : 101.6.158.37
Subnet Mask : 255.255.255.0
Gateway : 101.6.158.1
Total IP Address : 15
Ending IP Address : 101.6.158.51
Admin Network Type : 1/10Gbit Copper Non-Bonded

Server Components

Rack U Location	Component	NET0 Name	NET0 IP Address	ILOM Name	ILOM IP Address	VLAN ID
X6-2 Quarter Rack HC 8TB						
17	Database Server	dexadb02	101.6.158.38	dexadb02-ilom	101.6.158.43	
16	Database Server	dexadb01	101.6.158.37	dexadb01-ilom	101.6.158.42	
6	Storage Server	dexacel03	101.6.158.41	dexacel03-ilom	101.6.158.46	
4	Storage Server	dexacel02	101.6.158.40	dexacel02-ilom	101.6.158.45	
2	Storage Server	dexacel01	101.6.158.39	dexacel01-ilom	101.6.158.44	

Switch and Power Components

Rack U Location	Component	Admin Name	Admin IP Address
X6-2 Quarter Rack HC 8TB			
22	Leaf Switch	dexasw-ibb01	101.6.158.49
21	Admin Switch	dexasw-adm01	101.6.158.47
20	Leaf Switch	dexasw-iba01	101.6.158.48
0	PDU Switch	dexasw-pdua01	101.6.158.50
0	PDU Switch	dexasw-pdub01	101.6.158.51

Client Access Network Configuration

All IP addresses and corresponding host names should be registered in DNS
The Client network defines the default gateway for all Compute nodes

Starting IP Address : 101.6.132.37
Subnet Mask : 255.255.255.0
Gateway : 101.6.132.1
Total IP Addresses : 7
Ending IP Addresses : 101.6.132.43
Client Network Type : 10 Gbit optical Bonded

LACP : Disabled
BONDING_OPTS="mode=active-backup miiimon=100 downdelay=2000 updelay=5000 num_grat_arp=100"

Rack U Location	Component	Client Name	Client IP Address	VIP Name	VIP IP Address	VLAN ID
X6-2 Quarter Rack HC 8TB						
17	Database Server	dexadb02	101.6.132.39	dexadb02-vip	101.6.132.40	
16	Database Server	dexadb01	101.6.132.37	dexadb01-vip	101.6.132.38	

Private InfiniBand Network Configuration

Starting IP Address : 192.168.10.10
Subnet Mask : 255.255.255.0
Total IP Addresses : 10
Ending IP Address : 192.168.10.19

Rack U Location	Component	BONDIB0 Name	BONDIB0 IP Address
X6-2 Quarter Rack HC 8TB			
17	Database Server	dexadb02-priv1	192.168.10.12
		dexadb02-priv2	192.168.10.13
16	Database Server	dexadb01-priv1	192.168.10.10
		dexadb01-priv2	192.168.10.11
6	Storage Server	dexacel03-priv1	192.168.10.18
		dexacel03-priv2	192.168.10.19
4	Storage Server	dexacel02-priv1	192.168.10.16
		dexacel02-priv2	192.168.10.17
2	Storage Server	dexacel01-priv1	192.168.10.14
		dexacel01-priv2	192.168.10.15

Exadata Configuration Diskgroup Sizing X6-2 Published

EXADATA TYPICAL RACK CONFIGURATIONS²

Rack Size	Database Servers and Cores	Storage Servers and Cores	HC Storage Capacity (raw)	EF Storage Capacity (raw)
Eighth Rack³	2x servers, 44 cores	3x servers, 30 cores for SQL offload	144 TB disk, 19.2 TB flash	38.4 TB flash
Quarter Rack	2x servers, 88 cores	3x servers, 60 cores for SQL offload	288 TB disk, 38.4 TB flash	76.8 TB flash
Half Rack	4x servers, 176 cores	7x servers, 140 cores for SQL offload	672 TB disk, 89.6 TB flash	179.2 TB flash
Full Rack	8x servers, 352 cores	14x servers, 280 cores for SQL offload	1,344 TB disk, 179.2 TB flash	358.4 TB flash

² Each rack is 42 RU (Rack Units) in height, has 2x redundant Power Distribution Units (PDUs), 2x 36-port QDR (40 Gb/s) InfiniBand switches and 1x 48-port Cisco Ethernet switch for administration. Included Spare Parts Kit Contains:

- 1 x 3.2 TB NVMe PCI Flash card and 1 x 8 TB High Capacity disk, or
- 1 x 3.2 TB NVMe PCI flash drive

³ Eighth Rack is the minimum Exadata configuration. Eighth Rack compute servers have half the cores enabled. Eighth Rack EF storage servers have half the cores and flash drives enabled. Eighth Rack HC storage servers have half the cores enabled and half the disks and flash cards removed.

EXADATA TYPICAL RACK CONFIGURATIONS: COMBINED METRICS (HC & EF)

Combined Metrics	Data Capacity (Usable) – Normal Redundancy ⁶		Data Capacity (Usable) – High Redundancy ⁶
	HC	EF	
Full Rack	HC	508 TB	399 TB
	EF	130 TB	102 TB
Half Rack	HC	254 TB	199 TB
	EF	65 TB	51 TB
Quarter Rack	HC	109 TB	85 TB
	EF	28 TB	22 TB
Eighth Rack	HC	54 TB	43 TB
	EF	14 TB	11 TB

Exadata Configuration Diskgroup Sizing Considerations

- Redundancy
- Layout
 - 80/20
 - 60/40

Disk Group Details

Diskgroup Layout : Legacy 72%18%10% Legacy 34%56%10% Create Sparse DiskGroup

DBFS DiskGroup : DBFS HIGH Size : default

DATA DiskGroup : DATA1 HIGH Size : 72%

RECO DiskGroup : RECO1 NORMAL Size : 18%

Sparse DiskGroup : SPARSE1 NORMAL Size : 10%

If this is a **critical production database**, Oracle recommends configuring the **DATA** diskgroup with **HIGH** redundancy

Disk Group:

Disk Group	Redundancy	Volume Size	OCR/Vote	Quorum
DBFS	HIGH	291G	Yes	Yes
DATA1	HIGH	69468G (72%)	No	No
SPARSE1	NORMAL	9648G (10%)	No	No
RECO1	NORMAL	17352G (18%)	No	No

Exadata Configuration Diskgroup Sizing & Cluster Summary

Clusters Information

Cluster:c0_clusterHome

Cluster Information:

Version 18.2.0.0.180417
Name cluster-dal
Home /u01/app/18.0.0.0/grid
Inventory Location /u01/app/oraInventory
Base Dir /u01/app/grid
Client Domain genpt.net
Backup Location external
Backup Domain genpt.net
Client Scan Name dexa-scan
Client Scan Port 1521
101.6.132.41
Client Scan IPs 101.6.132.42
101.6.132.43
Client Gateway 101.6.132.1
Client Subnet Mask 255.255.255.0
WriteBack Flash Cache Enabled

Cluster Owner and Groups

Cluster User	Cluster Id	Groups	
		Name	Id
grid	100	oinstall	1010
		asmdba	1050
		asmoper	1060
		asmadmin	1070

Database:

Version 12.1.0.2.170418
Name dbm01
Database Home /u01/app/oracle/product/12.1.0.2/dbhome_1
Inventory Location /u01/app/oraInventory
Block Size 8192
Database Template OLTP
Base Dir /u01/app/oracle
[dexadbadm01.noname.net](#)
Database Machines [dexadbadm02.noname.net](#)

Database Owner and Groups

User	ID	Groups	
		Name	Id
oracle	1030	oinstall	1010
		asmdba	1050
		dba	1005
		racoper	1040

Disk Group:

Disk Group	Redundancy	Volume Size	OCR/Vote	Quorum
DBFS	HIGH	291G	Yes	Yes
DATA1	HIGH	69468G (72%)	No	No
SPARSE1	NORMAL	9648G (10%)	No	No
RECO1	NORMAL	17352G (18%)	No	No

Disk Group Storage

Storage	# of Disks	Total Disk Space(G)
dexaceladm01.noname.net	12	96000
dexaceladm02.noname.net	12	96000
dexaceladm03.noname.net	12	96000

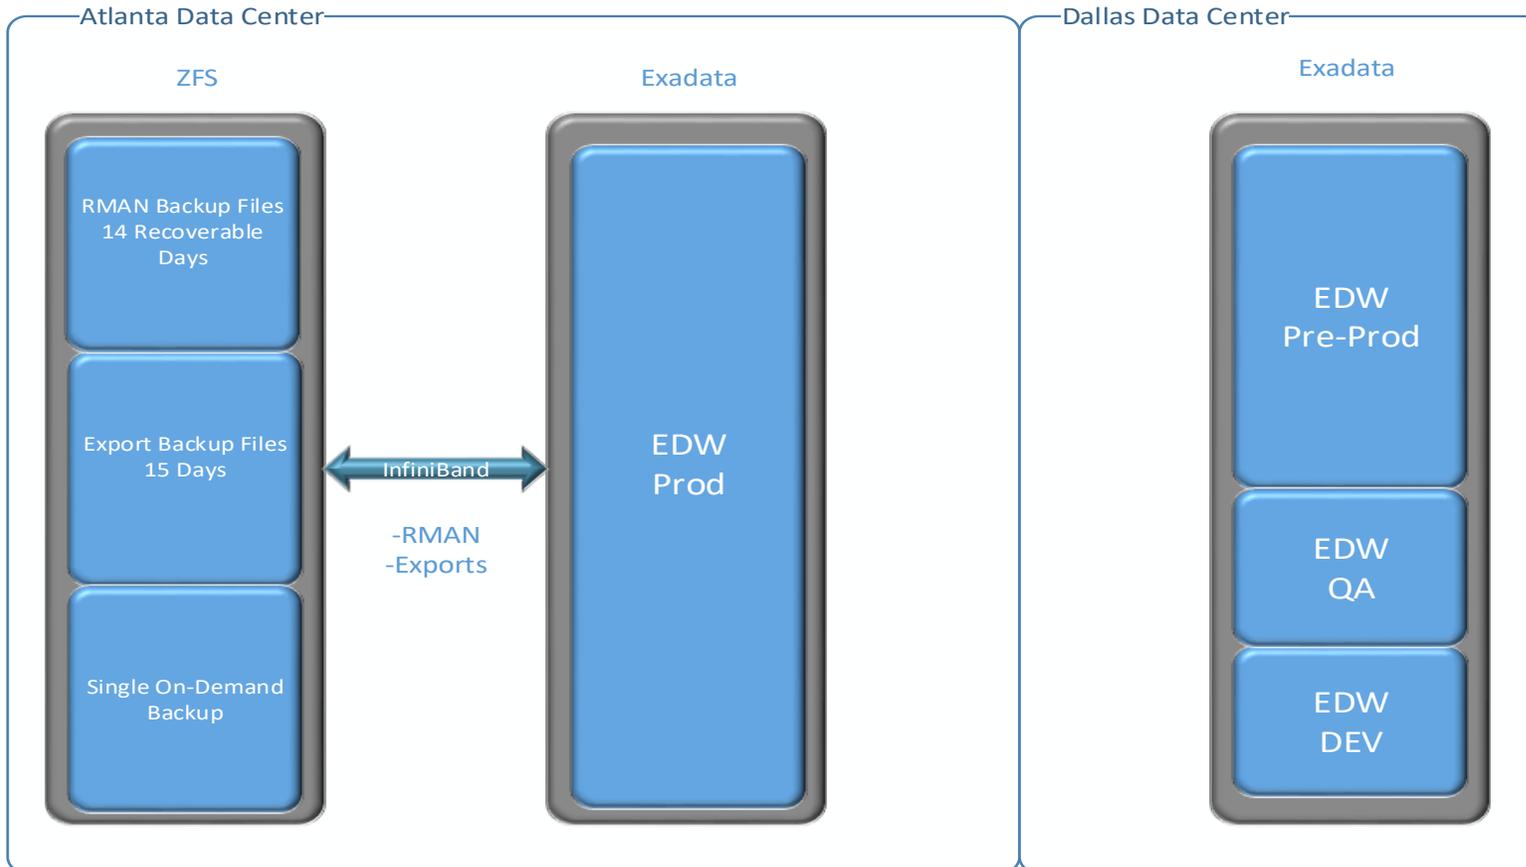
Exadata Configuration Installation

BIAS

```
Marks-iMac:macosx-x64 mdsaltzman$ ./install.sh -cf ExadataConfigurations/Genuine_Parts_Company-dexa.xml -l
  Initializing
  |||||
  1. Validate Configuration File
  2. Setup Capacity-on-Demand
  3. Setup Required Files
  4. Create Users
  5. Setup Cell Connectivity
  6. Verify Infiniband
  7. Calibrate Cells
  8. Create Cell Disks
  9. Create Grid Disks
 10. Install Cluster Software
 11. Initialize Cluster Software
 12. Install Database Software
 13. Relink Database with RDS
 14. Create ASM Diskgroups
 15. Create Databases
 16. Apply Security Fixes
 17. Install Exachk
 18. Create Installation Summary
 19. Resecure Machine
```

State of Source X2 Exadata

BIAS



ATLANTA

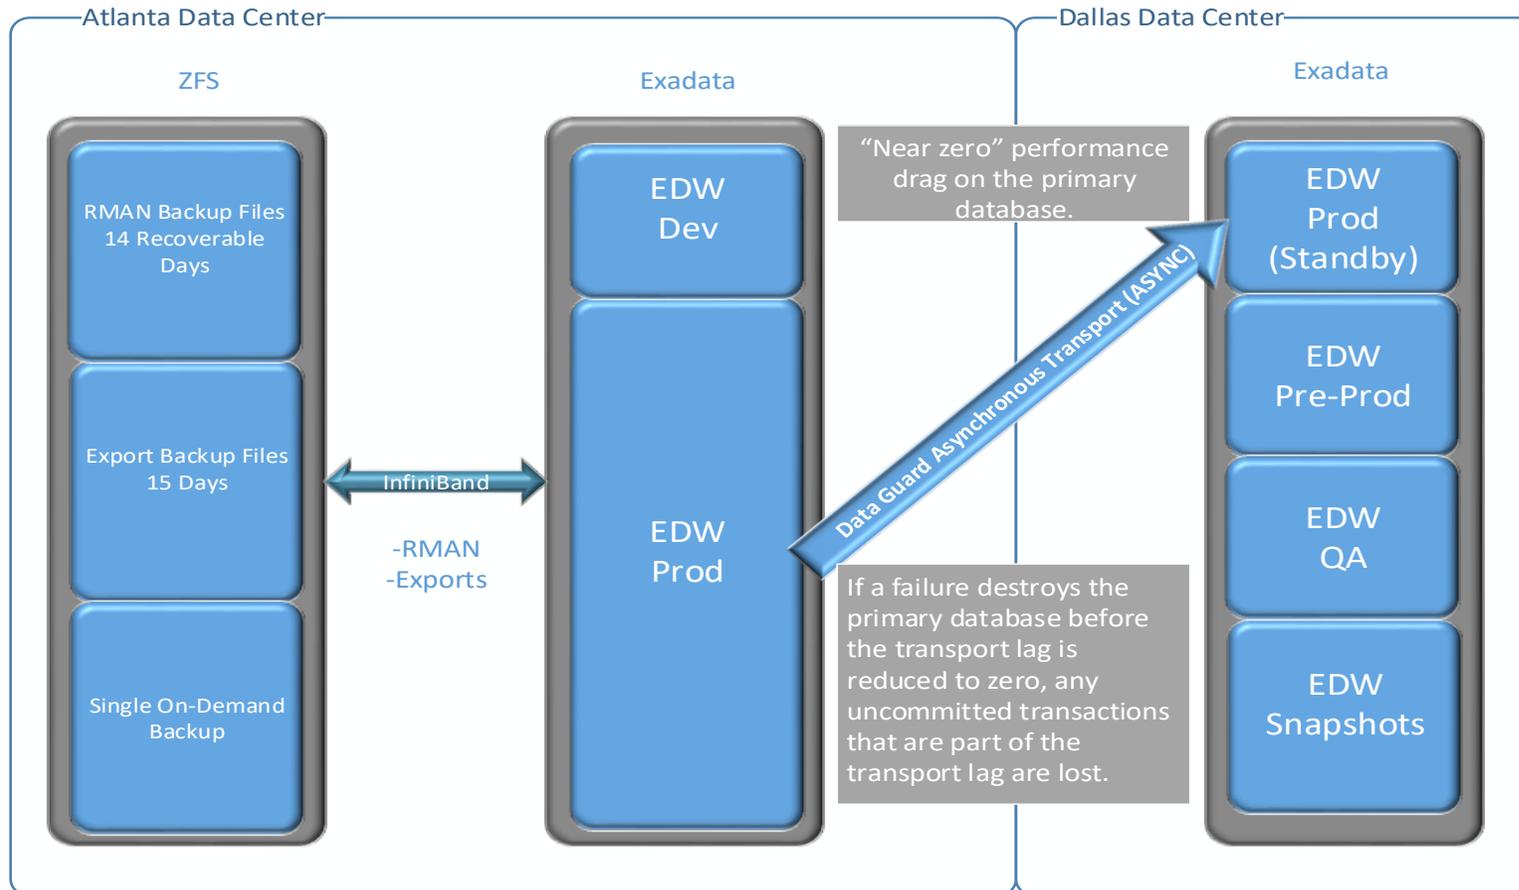
- 2 Compute Nodes
 - 24 cores
 - 192 GB RAM
- 3 Storage Cells
 - 36 cores
 - 1.1 TB Flash Cache
 - 10 TB usable storage

DALLAS

- 2 Compute Nodes
 - 24 cores
 - 192 GB RAM
- 3 Storage Cells
 - 36 cores
 - 1.1 TB Flash Cache
 - 10 TB usable storage

Proposed State of Target X6 Exadata

BIAS



ATLANTA

- 2 Compute Nodes
 - 20 cores
 - 512 GB RAM
- 3 Storage Cells
 - 60 cores
 - 38 TB Flash Cache
 - 144 TB usable storage

DALLAS

- 2 Compute Nodes
 - 16 cores
 - 512 GB RAM
- 3 Storage Cells
 - 60 cores
 - 38 TB Flash Cache
 - 144 TB usable storage

Migration Development Expdp/Impdp

- Create a new ACFS(/data)=300G and migrate edwdevpj data from X2(tx0001xdadb01.genpt.net) to here
- Please create dummy DB “EDWDEV” using DB character set/initialization parameters of EDWDEVPI(X2)
- Restore(exp/imp) database structures, tablespaces, permissions from production EDW(X2) to the EDWDEV database(X6)
- Restore data(exp/imp) ETL_USER schema from production EDW(X2) to the EDWDEV database(X6)
- Restore metadata(exp/imp) of EDW from production EDW(X2) to the EDWDEV database(X6)
- Restore(exp/imp) all production dimensions from production EDW(X2) to the EDWDEV database(X6) (any object that starts with D_)
- Restore(exp/imp) 1 month of data (most recent full month) from requested partitioned tables from X2 EDW PROD data to X6 EDWDEV
- Create DB link from EDWDEV DB(X6) to production EDW(X2)
- Restore(exp/imp) from X2 EDW PROD data to X6 EDWDEV 1 day of data (most recent) from the requested partitioned tables
- Shut down EDWDEVPI database on old X2 exadata servers
- Ensure that OEM monitoring is enabled for EDWDEV database(X6)
- Ensure that Oracle Platinum monitoring & metalink is enabled for new DB

- Disable scheduled jobs from OEM (Archive & Lev0/1 jobs)
- Kick off one-off RMAN backup in nohup mode
 - backup blocks all as compressed backupset incremental level 0 database tag 'edw_edwbkp' plus archive log tag 'edw_edwbkp' delete all input;
- Once backup completes, rename the backup destination directories
 - would not get deleted by any retention policy
- Enable the regular backups scripts from OEM
- Copy the backup pieces to destination server
- Restore the controlfile, update (unset) the unwanted rman parameters and clone the database using the backup pieces

```
CONFIGURE CHANNEL 1 DEVICE TYPE DISK FORMAT '/zfssa/edw/backup1/%U' CONNECT '*';  
CONFIGURE CHANNEL 2 DEVICE TYPE DISK FORMAT '/zfssa/edw/backup3/%U' CONNECT '*';  
CONFIGURE CHANNEL 3 DEVICE TYPE DISK FORMAT '/zfssa/edw/backup2/%U' CONNECT '*';  
CONFIGURE CHANNEL 4 DEVICE TYPE DISK FORMAT '/zfssa/edw/backup4/%U' CONNECT '*';
```

```
CONFIGURE SNAPSHOT CONTROLFILE NAME TO '+RECO/edw/snapcf_edw.f';
```

- Crosscheck backupset / archivelog all. Delete expired / obsolete backupset / archivelogs.
- Catalog the backup pieces to control file
- Run “list backup” and validate if there is any piece with elapsed time 00:00:00. If so, we need to re-catalog that piece.
- Set the db_name as source (EDW) and db_unique_name, instance_name as target name (ex: EDWQA) so that the files are created with correct directory structure in ASM.
- Ensure db_file_name_convert and log_file_name_convert parameters are set appropriately before kicking off RMAN restore
 - *.db_file_name_convert='+DATA','+DATA_DAL','+RECO','+RECO_DAL'
 - *.log_file_name_convert='+DATA','+DATA_DAL','+RECO','+RECO_DAL'
- Ensure to validate the total no of datafiles/tempfiles is accurate in “SET NEWNAME FOR datafile” and “SET NEWNAME FOR TEMPFILE”

Migration Production Using Switchover Method

- Disable scheduled jobs from OEM (Archive & Lev0/1 jobs)
- Gather Server Details and Database Details of Old Exadata X2 PRODUCTION and Exadata X6 (where Standby needs to be built)
 - DB Server, DB Unique Name, DB Home, DB User, DB Disk group
- Connecting ZFS storage to X6 servers
- Clone ORACLE HOME from Exadata X2 PRODUCTION(Source) to Exadata X6 (where Standby needs to be built)
- Create pfile for Standby from Primary & Take backup of Production Database using standby control file
- Copy tnsnames.ora entries for Standby Build on all nodes(2 nodes of primary on Old X2, 2 nodes of standby on new X6)
- Ensure Primary database Archive mode and standby config on primary(To be done once Standby Restore is completed)
- Create password file on Primary(To be done once restore is completed)
- Start the standby in nomount and restore control file.
- Initiate database restore on Exadata X6 and monitor the run until completion
- Start the MRP on Standby and Validate DG Sync
- RAC Conversion on new Standby on X6
- Switchover and bringing up Primary on X6.
- Enable all scheduled OEM and crontab jobs on Primary on X6.
- DR Build using above method on DR site

Issue 1: This is a known Oracle Bug where RMAN creates corrupted free blocks during recovery and causes related errors. There is no fix to this issue and work around is to use “blocks all” in RMAN backup script (e.g. backup **blocks all** as compressed backupset incremental level 0 database tag ‘edw_edwbkp’ plus archive log tag ‘edw_edwbkp’ delete all input) to avoid this issue.

ORA-00283: recovery session canceled due to errors

ORA-00600: internal error code, arguments: [3020], [219], [709377], [919261953], [], [], [], [], [], [], []

ORA-10567: Redo is inconsistent with data block (file# 219, block# 709377, file offset is 1516249088 bytes)

ORA-10564: tablespace ODS

ORA-01110: data file 219: '+DATA_DAL/DWPREPRD/DATAFILE/ods.501.948545309'

ORA-10560: block type '58'

Issue 2: Please disable block tracking after DB is mounted and before kicking off DB restore/recovery to avoid below error

ORA-19755: could not open change tracking file

ORA-19750: change tracking file: '+DATA/edw/changetracking/ctf.281.805121137'

ORA-17503: ksfopn:2 Failed to open file +DATA/edw/changetracking/ctf.281.805121137

Migration Issues



Issue 3: Some datafiles were getting restored with old SCN .The pieces need to be catalogued again and then initiate a restore of datafile followed by a switch.

- Check the logs the datafile instead of restoring there would be something like “creating datafile **”
- Cancel the restore, do a cleanup.
- Catalog the pieces again.

Rman target /

List copy of datafile **; -- if there is any O/P

Delete copy of datafile **;

Run

```
{
```

```
Set newname for datafile ** to '+DATA';
```

```
Restore datafile **;
```

```
}
```

Switch datafile ** to copy;

SQL> select distinct to_char(CHECKPOINT_CHANGE#),CHECKPOINT_TIME from v\$datafile_header; -- It should show consistent SCN.

- Recover the database with set until time clause.

Migration Issues



Issue 4: RMAN ACTIVE DUPLICATE USING BACKUPSET IN 12C for PRODUCTION Migration

RMAN ACTIVE DUPLICATE USING BACKUPSET IN 12C(NEW FEATURE) (Doc ID 1987193.1)

- Network Bandwidth was not supportive

Issue 5: X6 Exadata Server gets rebooted on its own and cluster is not coming up and one the IB switches were not pingable

*/*Short Term Plan – Immediate*/*

reboot IB switches via ilom

reboot cells

reboot compute nodes

*/*Long Term Plan */*

Shutdown CRS & login to atlexasw-iba01

run “getmaster” -----→ Ensure atlexasw-iba01(non-problematic switch) is master and is running opensmd

run “service opensmd status” -----→ Ensure atlexasw-iba01(non-problematic switch) is master and is running opensmd

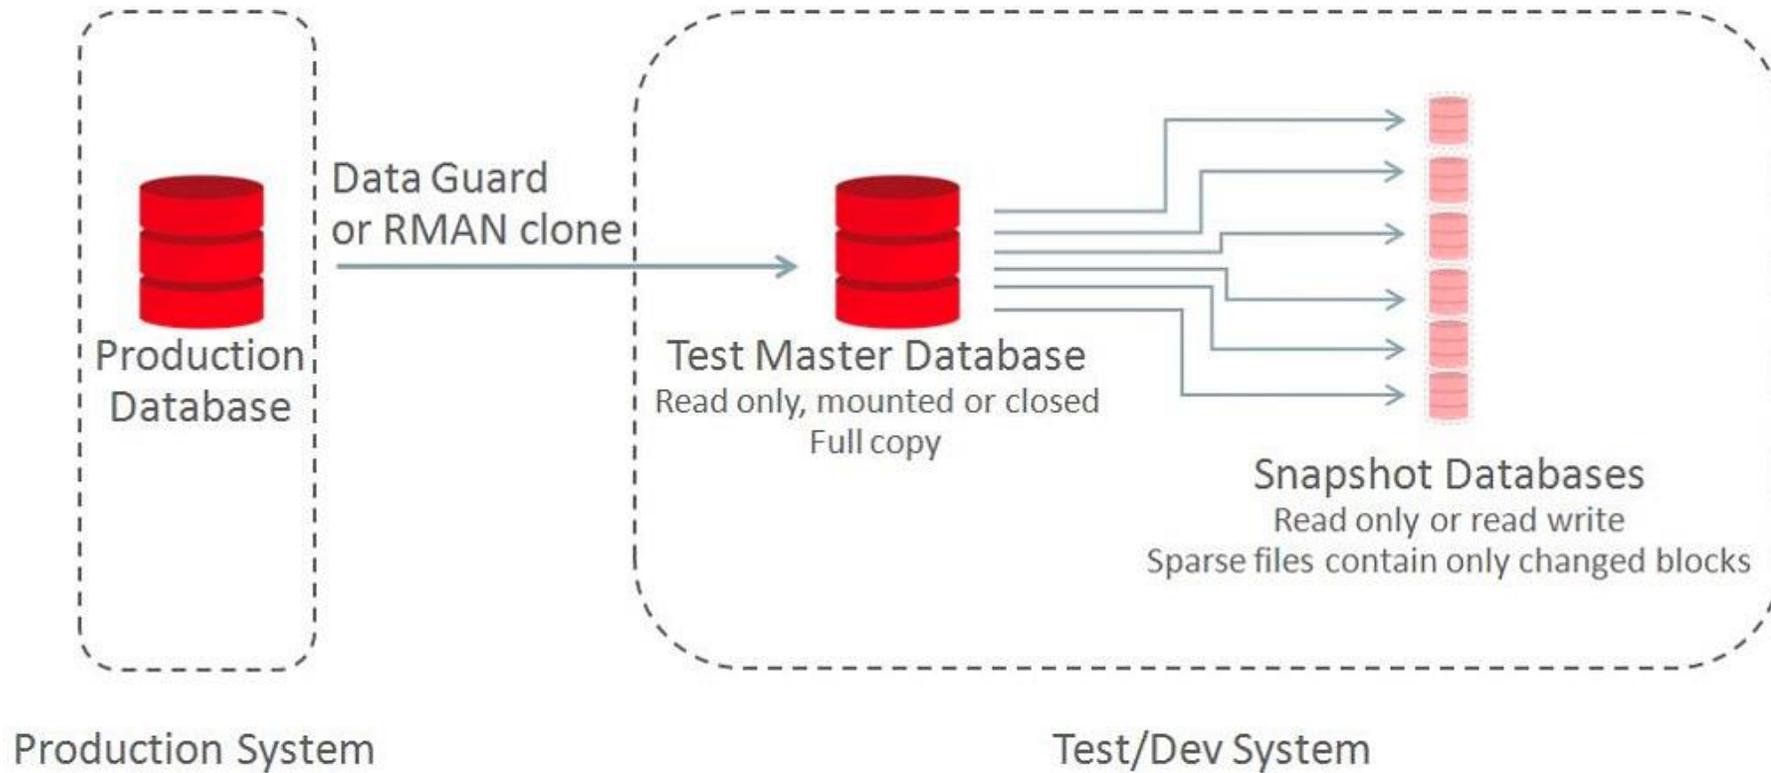
login to atlexasw-ibb01 and run disablesm

FE replace ibb0 switch as per SR action plan

login to atlexasw-ibb01 and run enablesm

Check ibswitches are communicating to all 2 compute nodes and 3 cell nodes using command “iblinkinfo” & Start CRS

Snap Clone Exadata Snapshots



Snap Clone

Pre-requisites



- Storage servers must be Exadata X3 or later
- Exadata Storage Server software 12.1.2.1.0 or later for Exadata Storage and Database Servers

You cannot downgrade to an earlier version with sparse grid disks on a cell.

- Oracle Grid Infrastructure 12.1.0.2.0 BP5 or later

The ASM disk group that contains the sparse ASM grid disks must have both COMPATIBLE.RDBMS and COMPATIBLE.ASM set to 12.1.0.2 or later.

The parent disk group can be 11.2 compatible.

- Oracle Database version 12.1.0.2.0 BP5 or later

The parent database and the snapshot database must be 12.1.0.2 compatible.

- The data files for the snapshot database and the parent database must be on the same ASM cluster.
- The db_block_size must be at least 4K and be a multiple of 4K.

Snap Clone SPARSE Diskgroup Requirements



A sparse disk group has the following attributes:

- compatible.asm must be set to 12.1.0.2 or higher.
- compatible.rdbms must be set to 12.1.0.2 or higher.
- cell.sparse_dg must be set to "allsparse".

This is a new attribute that identifies the disk group to ASM as being made up of sparse grid disks.

- appliance.mode must be set to true.
- A sparse disk group uses 16 times the size of an extent. For 4M allocation unit (AU), each extent is 64M.
- A sparse disk group uses Virtually Allocated Metadata.

For example, the following SQL command creates a sparse disk group:

```
SQL> create diskgroup SPARSE normal redundancy disk 'o/*/SPARSE_*' attribute 'compatible.asm' = '12.1.0.2', 'compatible.rdbms' = '12.1.0.2',  
'cell.smart_scan_capable' = 'true', 'cell.sparse_dg' = 'allsparse', 'au_size' = '4M';
```

Snap Clone TESTMASTER

- Enable access control on the disk group that contains the test master's data files. The disk group must be on Exadata.

```
ALTER DISKGROUP DATA3 SET ATTRIBUTE 'ACCESS_CONTROL.ENABLED' = 'TRUE';
```

Setting the Ownership of the Test Master Data Files

- Add an operating system user as owner of the disk group.

```
ALTER DISKGROUP DATA3 ADD USER 'oradebs';
```

- Generate a script called set_owner.sql to set the owner of the test master's data files.



- Remove extra lines in set_owner.sql.

```
sed -i '/SQL/d' set_owner.sql
```

- Run the script in the ASM instance.

```
SQL> @set_owner
```

- Remove write permissions on all the data files to help prevent accidental overwrite.

SQL commands in ASM instance only allow to set file permissions to read only. You cannot remove write permissions in SQL.

```
SQL> ALTER DISKGROUP DATA3 set permission owner=read ONLY, group=read ONLY,other=none for file 'FILENAME';
```

This allows snapshots to be created and owned by users other than the owner of the base files.

Snap Clone

Creating the Snapshot Databases



Step 1. In the test master database, create a sample control file script to use for your Exadata snapshot databases by backing up the existing control file to trace.

```
SQL> alter database backup controlfile to trace;
```

Step 2. In the test master database, determine the existing file names for the rename step

Log into SQL*Plus as sysdba and run the following:

```
set newpage 0
set linesize 999
set pagesize 0
set feedback off
set heading off
set echo off
set space 0
set tab off
set trimspool on
spool debs_rename_files.sql
select 'EXECUTE dbms_dnfs.clonedb_renamefile('||''''||name||''''||','||''''||replace(replace(replace(name,',' ,'_'),'MSTR','DEBS'),'DATA3','SPARSEC1')||''''||')' from v$datafile;
exit
```

Step 3: Shutdown TESTMASTER database

Step 4: Create an init.ora file for the target Exadata snapshot database.

Step 5. Edit the trace file generated in step 1 to create a control file for the Exadata snapshot database and create a .sql file called crt_ctlfile.sql that you will run later in step 7.

The control file should be created with Exadata snapshot database name, new log file names(any disk group that has enough space, but they should not be created in the sparse Oracle ASM disk group), and the test master's data file names.

Snap Clone

Creating the Snapshot Databases



Step 6: Start a database instance pointing to the Exadata snapshot database init.ora file using the following commands:

```
SQL> sqlplus / as sysdba
```

```
SQL> startup nomount pfile=initDEBS.ora
```

Step 7: Create the Exadata snapshot control file using the script created in step 5. In the following example the script is named crt_ctlfile.sql.

```
SQL> @crt_ctlfile
```

Step 8: Run the script you modified in step 2. All the files must be renamed prior to opening the Exadata snapshot database.

Connect using SQL*Plus as sysdba to the Exadata snapshot database and run the following:

```
SQL> @debs_rename_files
```

This script modifies the permissions of the test master database files in ASM, marking them as READONLY.

The dbms_dnfs.clonedb_renamefile procedure, which is called by rename_files.sql, sets up the parent-child relationship between the test master database and the snapshot database, and renames the filenames in the snapshot database's control file.

Step 9. Open the Exadata snapshot database with the RESETLOGS option: SQL> ALTER DATABASE OPEN RESETLOGS;

Step 10. Confirm that the Exadata snapshot files are child files of the test master database. Connect using SQL*Plus as sysasm to the Exadata snapshot, and run the following

```
SQL> select filenumber num, CLONEFILENAME child, SNAPSHOTFILENAME parent from x$ksfdsscloneinfo;
```

Step 11. Log in using SQL*Plus to the Exadata snapshot database, and add temp files to the TEMP tablespace.

This will be a full size temp file, not sparse.

BIAS

Mark Saltzman

Senior Director
Database Services
BIAS Corporation

Mark.Saltzman@biascorp.com

☎ 954-682-4733

Heema Satapathy

Senior Principle Consultant &
Service Delivery Manager
BIAS Corporation

Heema.Satapathy@biascorp.com

☎ 1-770-685-6301



ORACLE
ACE Associate