

Exadata Deployment Life Cycle

BY UMAIR MANSOOB



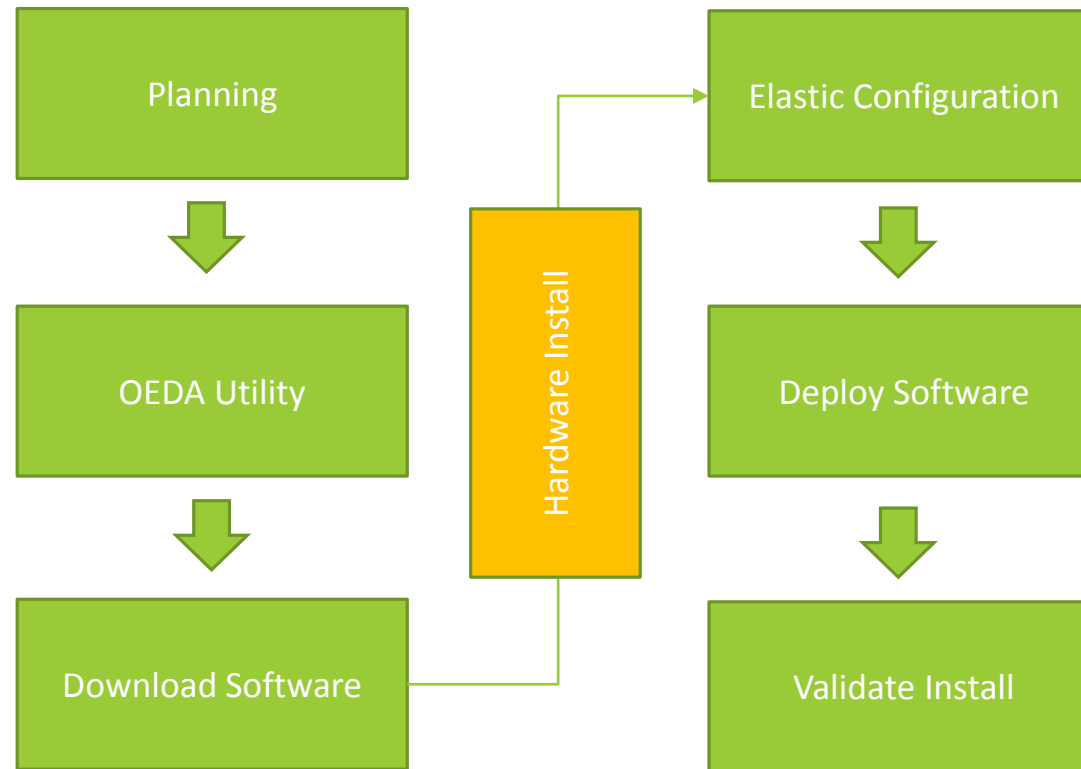
Agenda

- Introduction
- Planning
- OEDA Utility
- Hardware Install
- Download Software
- Elastic Configuration
- Software Install
- Validate Install

Who Am I

- Oracle Certified Administrator from Oracle 7 – 12c
- Exadata Certified Implementation Specialist since 2011
- Oracle Database Performance Tuning Certified Expert
- Oracle Business Intelligence Foundation Suite 11g Certified Implementation Specialist
- Oracle Database Data Warehousing Certified Implementation Specialist
- Multiple Exadata Implementations / POC's for large financial organizations
- Migrate / Upgrade databases between various versions of Oracle
- Capacity Planning for Oracle Engineered Systems
- Database Consolidation to Exadata / 12c Platform
- Architect Databases for OLTP and OLAP applications
- **Not an Oracle Employee or Nor I represent Oracle in any way**

Overview



Planning

Capacity Planning

Deployment Decisions

Licensing Consideration

Capacity Planning

- It is important to understand that Exadata offloads a lot of processes to storage nodes
- Capacity planning for Exadata is a little different from traditional hardware but basics will still apply
- You should gather information about CPU, IOPs and Physical Memory of an existing system and use them to size your Exadata machine

Deployment Consideration

- Exadata deployment decision can have lasting impact on your hosted databases
- Oracle recommends using High redundancy level for Exadata machine
- It is also important to understand that virtualizing Exadata machine comes with administrative overhead
- You can achieve network isolation using two options, private Vlan and InfiniBand partitioning
- Exadata machine comes with two storage disk options, Extreme Flash IO and high capacity.

Licensing Consideration

- As a customer, you have two options to control Exadata Software licensing cost
- You can use Capacity on Demand (CoD) option to disable a subset of the cores on Exadata database servers to reduce licensing requirements
- The maximum number of cores that can be disabled on an Exadata X5-2 Database Server is 60%
- OVM can be used to create virtual machines with specific number of cores to reduce licensing requirements,
- You might be required to buy Oracle advance Security option for compliance reasons especially, if your database contain credit card information
- Even though Exadata comes with its own free compression called Hybrid Column compression, it does not support OLTP operations

OEDA Utility

- Run OEDA configuration to Generate Configuration files
- You can run OEDA utility before the arrival of Exadata HW
- Work network and database administrators to gather following information
 - IP addresses
 - Machine names
 - DNS information
 - NTP information for the rack
- You can Download Latest OEDA Utility from following link
 - (<http://www.oracle.com/technetwork/database/exadata/oeda-download-2076737.html>)
- Oracle Exadata Deployment Assistant has two main phases
 - the customer provides naming, DNS, NTP, and networking details. This information is collected using the Oracle Exadata Deployment Assistant user interface. A configuration file is generated
 - the configuration file is pushed to the database node, and the validation and installation is completed






OEDA Download

Oracle Exadata Deployment Assistant Downloads

You must accept the [OTN License Agreement](#) to download this software.

Accept License Agreement | Decline License Agreement

Oracle Exadata Deployment Assistant

-  [OEDA for Microsoft Windows \(32/64-bit\) \(77 MB\) Aug 2016](#)
-  [OEDA for Linux x86-64 \(83 MB\) Aug 2016](#)
-  [OEDA for Mac OS X \(64-bit\) \(83 MB\) Aug 2016](#)
-  [OEDA for Oracle Solaris on SPARC \(64-bit\) \(111 MB\) Aug 2016](#)
-  [OEDA for Oracle Solaris on x86-64 \(64-bit\) \(107 MB\) Aug 2016](#)

OEDA – Welcome

Welcome

Select Oracle Exadata Database Machine if this is a **new** Exadata Deployment

Select Zero Data Loss Recovery Appliance if you are deploying a **ZDLRA**

Select Oracle SuperCluster if you are deploying a **SPARC based SuperCluster**

- Oracle Exadata Database Machine
- Zero Data Loss Recovery Appliance
- Oracle SuperCluster

OEDA – Customer Detail

Customer Details

Customer Name :

Application :

Network Domain Name :

Name Prefix :

Region :

Timezone :

DNS :

NTTP :

OEDA – HW Selection

Hardware Selection

Select Interconnected hardware to deploy

- X6-2
 - Full Rack
 - Half Rack
 - Quarter Rack
 - X6-2 Quarter Rack EF 3.2TB
 - X6-2 Quarter Rack HC 8TB
 - Eighth Rack
 - Elastic Rack
- X6-8
 - Full Rack
 - Elastic Rack
- X5-2
 - Full Rack
 - Half Rack
 - Quarter Rack
 - Eighth Rack
 - Elastic Rack

Add >

Remove <

Clear <<

This is your deployment

X6-2 Quarter Rack HC 8TB

OEDA – Rack Details

Rack Details

Rack 1 : X6-2 Quarter Rack HC 8TB

Compute Node count :

Storage Cell count :

Include a spine switch in this rack

OEDA – Define Network

Define Customer Networks

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 infinband network. Some customers have more than 2 subnets. In those cases Exadata can configure one of those additional subnets for 'backup', 'replication', 'dr' 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

Subnet 1

Name : **Admin** Bonded
Subnet Mask : 255.255.255.0 Non Bonded
Gateway : 193.28.223.254
Admin Network Format : 1/10 Gbit Copper Base-T 10 Gbit Optical

Subnet 2

Name : **Client** Bonded
Subnet Mask : 255.255.255.0 Non Bonded
Gateway : 193.68.93.254
Client Network Format : 1/10 Gbit Copper Base-T 10 Gbit Optical

Subnet 3

Name : **Private** Bonded
Subnet Mask : 255.255.252.0 Non Bonded
Private Network Format : InfiniBand

Subnet 4

Available Network : Backup Bonded
Subnet Mask : 255.255.255.0 Non Bonded
Gateway :
Backup Network Format : 1/10 Gbit Copper Base-T 10 Gbit Optical

OEDA – Administration Network

Administration Network

Starting IP Address for Pool :
[Valid network range : 193.28.223.1 - 193.28.223.253]

Pool Size :

Ending IP Address for Pool :

Is the default gateway for database servers

Defines the Hostname for the database servers

The pool should consist of consecutive IP addresses. If you cannot provide this then specific IP addresses can be modified at the end of the configuration process.

— Sample first host names —

Database Server Admin Name :	<input type="text" value="dm01dbadm01"/>	ILOM Name :	<input type="text" value="dm01dbadm01-ilom"/>
Storage Server Admin Name :	<input type="text" value="dm01celadm01"/>	ILOM Name :	<input type="text" value="dm01celadm01-ilom"/>

OEDA – Ethernet Network

Client Ethernet Network

Starting IP Address for Pool :
[Valid network range : 193.68.93.1 - 193.68.93.253]

Pool Size :

Ending IP Address for Pool :

Is the default gateway for database servers
 Defines the hostname for the database servers

The pool should consist of consecutive IP addresses. If you cannot provide this then specific IP addresses can be modified at the end of the configuration process.

Sample first database client names

Compute Client Name :	<input type="text" value="dm01db01"/>	VIP Name :	<input type="text" value="dm01db01-vip"/>
Client Scan name :	<input type="text" value="dm01-scan"/>		

OEDA – InfiniBand Network

InfiniBand Network

Starting IP Address for Pool :

Pool Size :

Ending IP Address for Pool :

Enable Active Bonding on Compute node Network

The pool should consist of consecutive IP addresses. If you cannot provide this, then specific IP addresses can be modified at the end of the configuration process.

[Sample first InfiniBand Name](#)

Compute Priv Name : Cell Priv Name :

OEDA – Identify Nodes

Identify Compute Node OS and Enable Capacity-on-Demand, if applicable.

Select the Operating System for the database servers

All Linux All Solaris All OVM

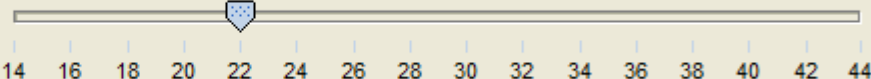
Solaris is not an available Operating System when Active Bonding is enabled on the InfiniBand Network.
If the Solaris OS is required then go back to the InfiniBand Network Page and de-select the Active Bonding Checkbox

- Rack : 1 Compute : 1 : Physical Linux
- Rack : 1 Compute : 2 : Physical Linux

Enable Capacity-on-Demand

OEDA – Enable Capacity on Demand

Enable Capacity-on-Demand



14 16 18 20 22 24 26 28 30 32 34 36 38 40 42 44

You have chosen to disable a portion of the cores on one or more database servers, a feature termed capacity-on-demand. To qualify for capacity-on-demand, within 3 months after installation this Exadata system must use either Platinum Services, Oracle Enterprise Manager Harvester (in disconnected or connected mode) or Oracle Configuration Manager (in connected mode).

OEDA – Mgmt & Private Network

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.

* 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.

Rack 1

Exadata X6-2 Compute Node 1		Rack 1 - Rack Location 16
Admin Name :	<input type="text" value="dm01dbadm01.my.company.com"/>	
Admin IP :	<input type="text" value="193.28.223.1"/>	
Ilom Name :	<input type="text" value="dm01dbadm01-ilom.my.company.com"/>	
Ilom IP :	<input type="text" value="193.28.223.6"/>	
Priv Name :	<input type="text" value="dm01db01-priv1.my.company.com"/>	Priv IP : <input type="text" value="192.168.10.1"/>

Exadata X6-2 Compute Node 2		Rack 1 - Rack Location 17
Admin Name :	<input type="text" value="dm01dbadm02.my.company.com"/>	
Admin IP :	<input type="text" value="193.28.223.2"/>	
Ilom Name :	<input type="text" value="dm01dbadm02-ilom.my.company.com"/>	
Ilom IP :	<input type="text" value="193.28.223.7"/>	
Priv Name :	<input type="text" value="dm01db02-priv1.my.company.com"/>	Priv IP : <input type="text" value="192.168.10.3"/>

OEDA – Define Clusters

Define Clusters

Number of Clusters to create :

Cluster 1

Cluster Name :

CELL dm01celadm01.my.company.com CELL dm01celadm02.my.company.com CELL dm01celadm03.my.company.com	<input type="button" value="Add >"/> <input type="button" value="All >>"/> <input type="button" value="Remove <"/> <input type="button" value="Clear <<"/>	dm01dbadm01.my.company.com dm01dbadm02.my.company.com dm01celadm01.my.company.com dm01celadm02.my.company.com dm01celadm03.my.company.com
--	---	---

OEDA – Cluster Configuration

Cluster 1

Cluster name : Physical Cluster

Prefix :

DNS :

NTP :

Domain Name :

Region : TimeZone :

Writeback Flash Cache

Users and Groups

Role Separated

Grid ASM Home OS User : ID : base :

ASM DBA Group : ID :

ASM Home Oper Group : ID :

ASM Home Admin Group : ID :

RDBMS Home OS User : ID : base :

RDBMS DBA Group : ID :

RDBMS Home Oper Group : ID :

OINSTALL Group : ID :

OEDA – Review Cluster Configuration

Review and Edit SCAN, Client, VIP and optional Backup Networks

* 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.

Cluster cluster-01

SCAN Name :

SCAN IP 1 :

SCAN IP 2 :

SCAN IP 3 :

Compute Node 1

Client Name : IP :

VIP Name : IP :

Compute Node 2

Client Name : IP :

VIP Name : IP :

OEDA – Grid Control Agent

Grid Control Agent

Enable Oracle Enterprise Manager Grid Control Agent

EM Home Base Location :

OMS Hostname :

OMS HTTPS Upload Port :

OEDA – Configuration Files

Your configuration files have been created

You can view the Installation template by clicking the following link

[Installation template](#)

Directory Location for Files :

C:\Users\lumansoob\Desktop\Software\Ocmd-16.238-OTN-windows-i586\windows-i586\ExadataConfigurations\

OEDA XML file : customer_name-dm01.xml

Network Check Script : customer_name-dm01-checkip.cmd

Installation Template : customer_name-dm01-InstallationTemplate.html

Preconf file(s) : customer_name-dm01-preconf_rack_0.csv
customer_name-dm01-platinum.csv

Download Software

- Please review following Notes 888828.1 for latest patch updates
- Download Required Software (GI , RDBMS , Latest Patches) bases on your Exadata Installation Template

Download the following files from MOS into the Oracle Exadata Deployment Assistant's WorkDir directory

```
linuxamd64_12102_database_1of2.zip  
linuxamd64_12102_database_2of2.zip  
linuxamd64_12102_grid_1of2.zip  
linuxamd64_12102_grid_2of2.zip  
p21923026_12102160419DBEngSysandDBIM_Linux-x86-64.zip  
p22674709_121020_Linux-x86-64.zip  
p22899531_121020_Linux-x86-64.zip  
p23200778_12102160419DBEngSysandDBIM_Linux-x86-64.zip  
p6880880_121010_Linux-x86-64.zip  
exachk.zip
```

Patch for bug 23200778 is required for all bundle patches. If it is not available on My Oracle Support for your database version, contact Oracle Support

Hardware Install

- Mostly Perform By Oracle
- Unpacking Oracle Exadata Rack
- Powering on the System
- Connecting to Network
- Checking Exadata Database & Storage Server
- Checking InfiniBand Network

Apply Elastic Configuration

- Download OEDA Utility and configuration on all DB nodes
- Run check IP to validate Network Configuration
- ApplyElasticConfig will Use your Exadata configuration file to assign new IP addresses to Machine
- All the nodes will be rebooted with new IP's
- Command Systac
 - `./applyElasticConfig.sh -cf customer_name-configFile.xml`
- Only to Apply Elastic Configuration from 1 node only

Don't forget to Run Checkip.sh

- Validate Network configuration
- Pingable IP addresses
- Not Pingable IP addresses
- Resolve IP Conflicts
- When to Run CheckIP
 - Before Applying Elastic Configuration
 - Before Deploying Exadata Software
 - After Applying Elastic Configuration
 - After Deploying Exadata Software
- Failure can cause Delay's
- `./checkip.sh -m {options}`

Elastic Configuration

```
[root@node8 linux-x64]# pwd
/opt/oracle.SupportTools/onecommand/linux-x64
[root@node8 linux-x64]# ./applyElasticConfig.sh -cf [REDACTED].xml
Applying Elastic Config...
Applying Elastic configuration...
Searching Subnet 172.16.2.x
6 live IPs in 172.16.2.x
Exadata node found 172.16.2.37
Configuring node : 172.16.2.46
Done Configuring node : 172.16.2.46
Configuring node : 172.16.2.40
Done Configuring node : 172.16.2.40
Configuring node : 172.16.2.38
Done Configuring node : 172.16.2.38
Configuring node : 172.16.2.37
Done Configuring node : 172.16.2.37
Configuring this node. Validate that it restarts with the right configuration
Configuring node : 172.16.2.44
|||||
Broadcast message from root@[REDACTED]dbadm01.[REDACTED].com
      (unknown) at 10:55 ...
I

The system is going down for reboot NOW!
Completed configuring node : 172.16.2.44
Completed elastic configuration
Completed Applying Elastic Config...
Ending applyElasticConfig
[root@node8 linux-x64]# █
```

Elastic Configuration - Before

```
root@node8:/opt/oracle.SupportTools
login as: root
root@172.16.2.44's password:
Last login: Thu Jul 28 14:35:15 2016
[root@node8 ~]# ibhosts
Ca      : 0x0010e00001905e28 ports 2 "node10 elasticNode 172.16.2.46,172.16.2.46 ETH0"
Ca      : 0x0010e000018ef2e8 ports 2 "node4 elasticNode 72.16.2.40,172.16.2.40 ETH0"
Ca      : 0x0010e000018fc4a8 ports 2 "node2 elasticNode 72.16.2.38,172.16.2.38 ETH0"
Ca      : 0x0010e00001902518 ports 2 "node1 elasticNode 72.16.2.37,172.16.2.37 ETH0"
Ca      : 0x0010e000010a2650 ports 2 "node8 elasticNode 72.16.2.44,172.16.2.44 ETH0"
```


Elastic Configuration - After

```
root@ [redacted] badm01:~  
login as: root  
root@172.18.216.79's password:  
Last login: Wed Aug 24 10:40:14 2016 from node8.my.company.com  
[root@ [redacted] badm01 ~]# ibhosts  
Ca      : 0x0010e00001905e28 ports 2 [redacted] dbadm02 S  
A-1"  
Ca      : 0x0010e000018ef2e8 ports 2 [redacted] celadm03 C  
HCA-1"  
Ca      : 0x0010e000018fc4a8 ports 2 [redacted] celadm01 C  
CA-1"  
Ca      : 0x0010e00001902518 ports 2 [redacted] celadm02 C  
CA-1"  
Ca      : 0x0010e000010a2650 ports 2 [redacted] dbadm01 S  
A-1"
```

Deploy Software

- There are total of 19 Steps at this point
- Need to have all the software and patches in (/u01/onecommand/linux-x64/WorkDir
- Need to Run it from 1 node only
- You can run all step together (Not Recommended)
- Command Syntax
 - `./install.sh -cf customer_name-configFile.xml -s {1-19}`

Deploy Software Steps

```
[root@dbadm01 linux-x64]# ./install.sh -cf [redacted].xml -l
Initializing

1. Validate Configuration File
2. Setup Required Files
3. Create Users
4. Setup Cell Connectivity
5. Verify Infiniband
6. Calibrate Cells
7. Create Cell Disks
8. Create Grid Disks
9. Configure Alerting
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
[root@dbadm01 linux-x64]#
```

Deploy Software Step 1 - Example

```
[root@dbadm01 linux-x64]# ./install.sh -cf [REDACTED].xml -s 1
Initializing
Executing Validate Configuration File
Validating cluster: [REDACTED]-clul
  Locating machines...
  Verifying operating systems...
  Validating cluster networks...
  Validating network connectivity...
  Validating NTP setup...
  Validating physical disks on storage cells...
  Validating users...
  Validating platinum...
  Validating switches...
  Validating quorum disk requirements...
  Checking disk reclaim status...
  Checking Disk Tests Status...
  Validating nodes for database readiness...
  Completed validation...
```

Deploy Software Step 19 - Example

```
[root@badm01 linux-x64]# ./install.sh -cf [REDACTED].xml -s 19
Initializing
Executing Resecure Machine
Stopping cluster [REDACTED]
Cluster [REDACTED] stopped...
Stopped clusterware on all nodes...
Making sure all cluster nodes are up...
Resecuring all nodes...
Waiting for machines to stop...
Waiting for machine [REDACTED] to be NOT pingable...
Waiting for machine [REDACTED] to be NOT pingable...
Waiting for machine [REDACTED] to be NOT pingable...
Waiting for machine [REDACTED] to be NOT pingable...
Waiting for machines [REDACTED] to be pingable...
Waiting for machine [REDACTED] to be pingable...
Waiting for machine [REDACTED] to be pingable...
Waiting for machine [REDACTED] to be pingable...
|||||
Broadcast message from root@[REDACTED].com
      (unknown) at 18:59 ...

The system is going down for reboot NOW!
Cell diag collection user CELLDIAG has been created on the following cell servers:
1. [REDACTED].com
2. [REDACTED].com
3. [REDACTED].com
Password of the CELLDIAG user has been reset to a random password. Please login to each of the cells as "celladmin " and change the password using the following command:
"cellcli -e ALTER USER CELLDIAG password=<newPassword>"
Successfully completed execution of step Resecure Machine [elapsed Time [Elapsed = 568283 ms [9.0 minutes] Wed Aug 24 18:59:56 CDT 2016]]
```

Validate Install

- Change root password for all the nodes including storage nodes
- Review Installation Summary Report
- Run Exachk Report
- Validate ASR configuration
- Configure Exadata Monitoring (OEM)

Thank You

Umair Mansoob

773-297-2061

umairmansoob@gmail.com

<http://blog.umairmansoob.com/>

