DELL POWEREDGE M420 AND ORACLE DATABASE 11G R2: A REFERENCE ARCHITECTURE

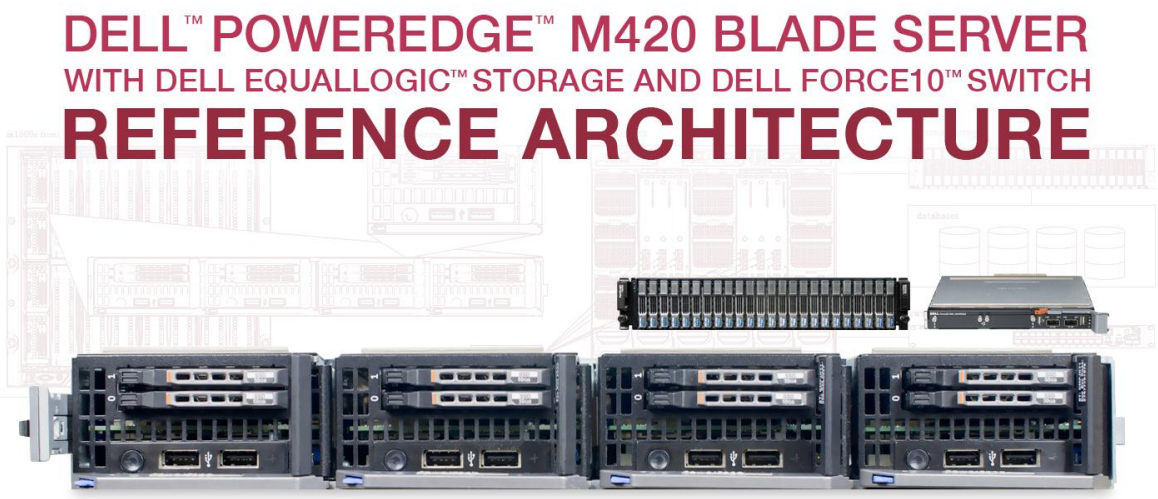

for clustering Oracle® Database 11g R2

Principled Technologies®

DELL

A Principled Technologies Reference Architecture commissioned by Dell Inc.

WHAT YOU WILL LEARN

- The features and advantages of the Dell PowerEdge M420
- How to install and configure your Dell PowerEdge M420 hardware
- How to install Oracle Database 11*g* R2 and Oracle RAC on the Dell PowerEdge M420
- How to configure Oracle Database 11*g* R2

TABLE OF CONTENTS

What you will learn **2**
Table of contents **2**
Executive summary **3**
About the test environment **3**
Features of the Dell PowerEdge M420 blade server 4
Features of the Dell Force10 MXL 10/40GbE blade switch 4
Features of the Dell EqualLogic PS6110XS 10GbE iSCSI array 5
About Oracle Real Application Clusters (RAC) 5
Supporting Hardware – PowerEdge M1000e and Fabric **6**
Configuring the storage 7
Configuring the server 9
Installing the OS and cluster prerequisites **10**
Installing Oracle Linux 6 10
Clustering prerequisites 11
Installing and configuring Oracle RAC **17**
Creating the cluster 17
Installing Oracle RAC 11*g* (Oracle Database) on the cluster 28
Summary **42**
Appendix A – Dell PowerEdge M420 system details **43**
About Principled Technologies **45**

EXECUTIVE SUMMARY

Data center space almost always comes at a premium. For that reason, it's important to consider not only performance and manageability, but also the amount of physical space your servers use when you are selecting a blade server to run your database applications. The ultra-dense, quarter-height Dell PowerEdge M420 blade server lets you fit twice as many servers in a single chassis than if you were using conventional half-height blades, without sacrificing performance or other features. Combined with high-performance storage and network options, such as the Dell EqualLogic PS6110XS storage array and Dell Force10MXL 10/40GbE switches, the M420 is able to provide excellent performance in a compact form-factor.

Dell designed the Dell PowerEdge M420 blade server specifically to meet the needs of these organizations that want to maximize their data center space. In this reference architecture, we look at some of the many benefits that the Dell PowerEdge M420 server can bring to data centers in combination with the EqualLogic PS6110XS and Force10MXL switches, and describe in detail the setup of a common clustering configuration: Oracle Real Application Clusters (Oracle RAC) with Oracle Database 11*g* R2 running on Oracle Linux 6.

ABOUT THE TEST ENVIRONMENT

To create this guide, we set up and tested an Oracle RAC with Oracle Database 11*g* R2 on Dell PowerEdge M420 servers. Figure 1 presents a concise configuration summary for the solution we tested.

Solution	Dell PowerEdge M420 solution
Servers	2 x Dell PowerEdge M420
Operating system	Oracle Enterprise Linux 6.2, x86_64
Database	Oracle Database 11*g* R2
Switches	2 x Dell Force10 MXL 10/40GbE
Storage	1 x Dell EqualLogic PS6110XS

Figure 1: General configuration summary for the solution we tested.

For more detail about the configuration of the servers we used, please see Appendix A.

Features of the Dell PowerEdge M420 blade server

The new Intel® Xeon® processor E5-2470-powered Dell PowerEdge M420 blade server (see Figure 2) sets a new standard for computing density in a blade environment, all using your existing Dell PowerEdge M1000e blade enclosure. Features include:

Figure 2: The Dell PowerEdge M420 blade server.

- **Double the computing power per rack unit.** Quarter-height blades capable of handling Intel Xeon processor E5-2400-series parts in a two-socket configuration, with up to eight cores per processor allow you to use your rack space to its maximum potential.
- **Cost-efficient upgrade.** Utilizing the same M1000e chassis as previous Dell blades, deploy the PowerEdge M420 using the power, cooling, and network infrastructure already in your data center.
- **Management.** The Dell OpenManage™ suite of tools, combined with the Integrated Dell Remote Access Controller (iDRAC7), give you maximum flexibility in administering your computing infrastructure.
- **Connectivity.** Snap-in FlexIO technology provides the backbone for a flexible, high-speed interconnect system. With options including 1 and 10 Gigabit Ethernet, Fibre Channel, InfiniBand®, and the new Dell Force10 MXL 10/40GbE blade switches, the M420 is able to interface with your network at maximum speed.
- **Power efficiency.** Dell PowerEdge M420 servers implement leading industry power-efficiency standards, and are designed to optimize performance while maintaining low power consumption.
- **Virtualization-ready.** Supporting Hyper-V™, Citrix® XenServer™, and VMware® vSphere®, the M420 contains two SD card slots for hypervisor, redundant hypervisor, and/or vFlash media support.

For details on the Dell PowerEdge M420 blade server configuration we used, see Appendix A.

Features of the Dell Force10 MXL 10/40GbE blade switch

The Dell Force10 MXL 10/40GbE blade switch brings new connectivity and performance options to blade computing environments with modular configuration options and standards-based 10/40GbE support. The Force10 MXL features:

- **Performance.** 10/40GbE connectivity minimizes latency and maximizes network throughput for data-intensive servers and applications, and Layer 3 routing guarantees fast traffic between isolated networks.

- **Flexibility.** Choose 1, 10, or 40GbE ports; FlexIO for plug-in QSFP+, SFP+; or 10GbE copper. Aggregate bandwidth on the FlexIO modules to support a maximum, bi-directional throughput of up to 160 Gigabits per second.
- **Extend and simplify.** The MXL is modular, and supports stacking for ease of management and higher performance. Support for converged switching allows iSCSI, NAS, Ethernet, and Fibre Channel traffic to all use the same physical hardware, potentially reducing additional management overhead and hardware costs.

Features of the Dell EqualLogic PS6110XS 10GbE iSCSI array

The Dell EqualLogic PS6110XS 10GbE iSCSI array combines the high-I/O strengths of solid-state storage and combines it with the raw capacity of traditional hard-disk drives in a single-chassis solution. Features include:

- **Agility.** Use storage that fits the demands of your server or application. Dell management software allows administrators to dynamically allocate storage on SSDs or HDDs, depending on the performance requirements for a specific situation.
- **Upgradable.** Start out using your existing 10GbE SFP+ modules, and build out to cost-effective 10GBASE-T as your budget allows. Move volumes between storage pools without downtime, and change hardware without disruption.
- **Robust.** Vertical port sharing keeps your bandwidth at full speed, even if a network port fails. Automatic load-balancing keeps data moving reliably to and from your physical or virtual server farm.
- **Management.** EqualLogic Host Software, Host Integration Tools, and EqualLogic SAN Headquarters (SANHQ) provide the tools administrators need to keep a high-level view of the storage system, while providing the flexibility to make granular changes as needed.

About Oracle Real Application Clusters (RAC)

Oracle Real Application Clusters (RAC) is a highly scalable cluster database solution for your business applications. Oracle RAC works to keep applications highly available by removing a single server as a point of failure for your infrastructure. By clustering servers, applications can continue to run on another server in the cluster should one fail. Scalable for infrastructures both small and large – up to 100 server nodes – Oracle RAC is included with Oracle Database 11*g* R2 software for ease of setup and configuration. For more information about Oracle RAC, visit

http://www.oracle.com/us/products/database/options/real-application-clusters/overview/index.html.

SUPPORTING HARDWARE – POWEREDGE M1000E AND FABRIC

The Dell PowerEdge M1000e blade chassis and its supported fabric interconnects are designed for large-scale database applications and application consolidation. Features of the PowerEdge M1000e include:

- **Management.** Reduces administrative demand by providing a secure centralized management interface for the chassis and blades within, using proven Web (SSL-encrypted) and CLI (SSH/Telnet) technologies.
- **Simplified configuration.** The Chassis Management Controller allows administrators to control up to nine enclosures and 144 server blades, including BIOS/firmware change management and updates, thermal monitoring, and power threshold configuration.
- **Flexible I/O.** Six interconnect sockets with the capability to support three fully-redundant fabrics, a passive midplane with more than 8Tbps in I/O bandwidth capacity, and FlexIO support provide a number of connectivity options for your servers.
- **Reliability and efficiency.** Six power supplies and nine fans, all hot-swappable, allowing for no-downtime maintenance of key chassis components. All components are tuned for maximum power efficiency to reduce data center power consumption.

All situations vary and you should consider your specific scenario when designing your topology. In our case, we used two Force 10 MXL 10/40GbE switches along with four Dell PowerConnect™ M8024-K switches for fabrics A, B, and C in our PowerEdge M1000e chassis.

For redundancy, we combined the Force10 MXL switches in a single stack to support our Oracle RAC intra-node traffic, and we assigned each NIC on Fabric A in our PowerEdge M420 nodes to one of two intra-node private networks. For Fabrics B and C, we stacked the PowerConnect M8024-K switches and bonded the third and fourth NICs in each node for redundancy and assigned them to the public network. At the same time, we used the EqualLogic Host Integration Tools for Linux, rather than OS's tools, to manage NICs three and four for multipath iSCSI traffic. This configuration allowed for higher redundancy and the highest level of intra-node and iSCSI throughput. Each networking situation will vary, so consult your network administrator and the network

hardware manuals for best practices on balancing performance and redundancy. Figure 3 shows our setup for this reference architecture.

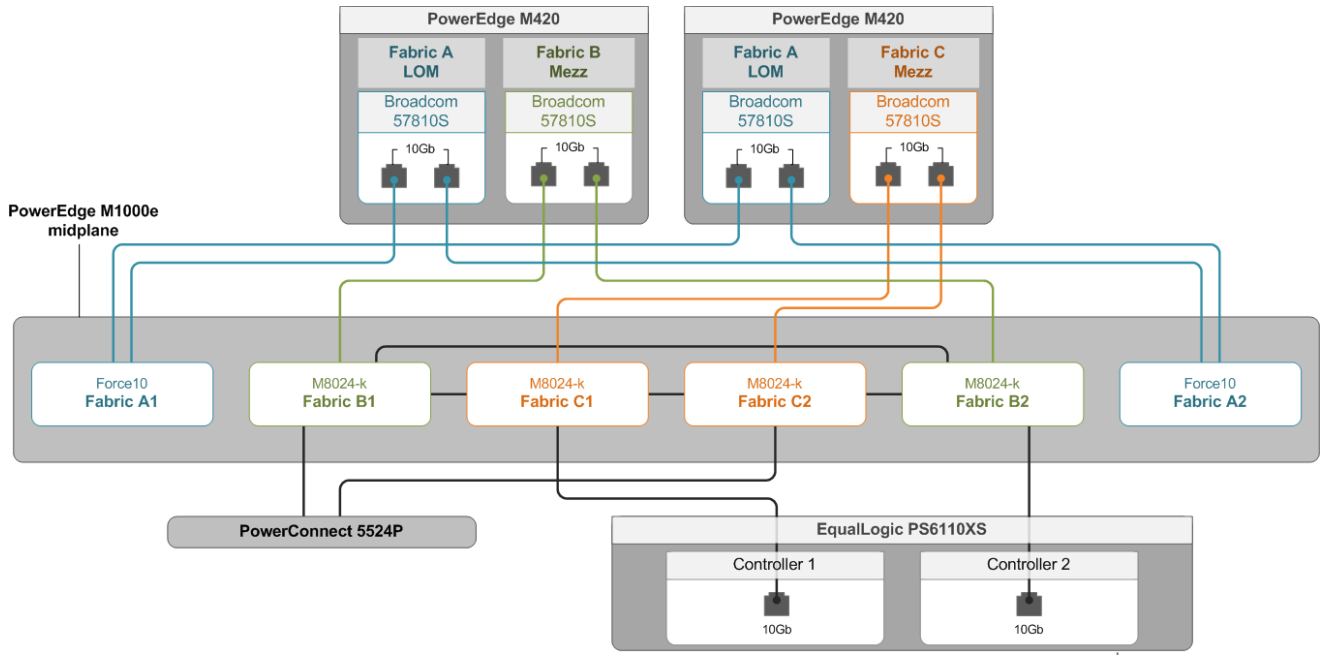

Figure 3: The topology we used for this reference architecture.

Configuring the storage

Storage considerations and best practices

The storage subsystem plays a vital role in any database system design. In our topology, we used a Dell EqualLogic PS6110XS as the main storage hardware for our Oracle Database 11*g* R2 environment. You should consider the following points when designing your storage for your environment:

- Always ensure that you have applied the latest firmware and updates to your storage subsystem and storage network hardware.
- If using iSCSI, consult the manufacturer's best practices for switch configuration, such as settings related to Jumbo Frames, Flow Control, Spanning tree, and so on.
- If using iSCSI, use VLANs to separate storage traffic when physically separate networks are unavailable.
- Redundancy is critical – always plan for and provide multiple paths for your hosts to communicate with your storage. Use multipath drivers from the storage hardware vendor to ensure optimal performance.

- Use some type of tiering – if your environment does not contain an EqualLogic PS6110XS or other storage with auto-tiering functionality, consider tiering your data manually, placing critical and volatile database files on faster storage, and placing data accessed less frequently on slower storage.
- Consider the workload – application workloads can vary greatly in data access patterns, and can therefore affect your system's performance. See the Dell white paper "Comprehending the Tradeoffs between Deploying Oracle Database on RAID 5 and RAID 10 Storage Configurations" for advice on RAID configurations.

Each situation and application design is different. Consider all layers of your hardware and software stack and their interoperability prior to establishing a storage design.

Configuring the Dell EqualLogic PS6110XS

Thanks to the unified management interface of Dell EqualLogic storage (see Figure 4), configuring the Dell EqualLogic PS6110XS is like configuring any other Dell EqualLogic array. To configure your storage, you should:

- Configure your group, if not already configured. You can perform this action via a serial connection to the first member array, using the setup command.
- Configure RAID levels on your arrays. In the Web interface, right-click the member, and choose Modify Raid Configuration.
- Create the volumes and configure authentication to those volumes. In the Web interface, click Volumes, and select Create volume. You can configure volume access from the Access tab.
- After configuration, use SAN HeadQuarters (SANHQ) to monitor your SAN for any issues.
- Monitor the Dell EqualLogic support site for any firmware or Host Integration Tool upgrades.

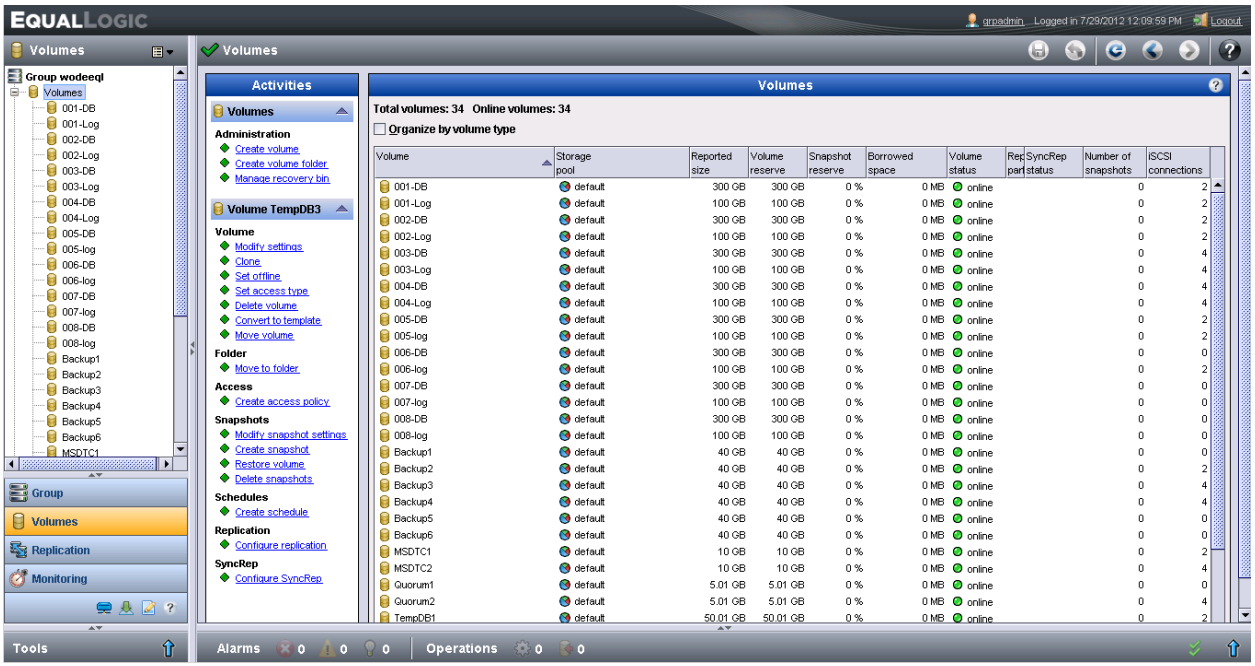

Figure 4: The Dell EqualLogic management Web interface.

Configuring the server

Dell PowerEdge M420 considerations and best practices

The Dell PowerEdge M420 blade server ships with the Intel Xeon processor E5-2400-series in a two-socket configuration and with up to 192GB RAM. When configuring your PowerEdge M420 hardware, consider the following:

- Most database workloads are RAM intensive. Configure your server with enough RAM to handle your database workloads.
- Core count options for the PowerEdge M420 are 4, 6, and 8. Processor cache options range from 10MB to 20MB cache. We recommended buying the processors with as many cores and as much cache as your budget allows, for maximum performance. This is especially true for CPU-intensive database application workloads that may thrive with many cores, faster frequencies, or more CPU cache.
- Ensure your PowerEdge M420 BIOS and firmware are at the latest levels released by Dell.
- Test your workload patterns and performance profiles using the Dell-supplied BIOS power options to maximize power efficiency.

INSTALLING THE OS AND CLUSTER PREREQUISITES

Installing Oracle Linux 6

For this reference architecture, we used Oracle Linux 6.2, the Red Hat Compatible Kernel version. You can download this software directly from Oracle via https://edelivery.oracle.com/linux/. Below we review the OS installation steps for installing Oracle Enterprise Linux on the Dell PowerEdge M420. Perform these actions for each server node.

Configuring the local disks and boot options

To install the operating system, first you must configure at least a single logical volume on the Dell PowerEdge M420. We chose to configure a RAID 1 mirrored pair of disks in a single volume for the operating system. Perform these steps on all servers.

1. Connect to the iDRAC management port of the individual PowerEdge M420.
2. Reboot the server, and press F10 to enter the Lifecycle Controller.
3. Select OS Deploy, and click Next.
4. Configure the local disks as a mirrored RAID 1 pair.
5. Connect the Oracle Linux 6 installation DVD to the server either through the server's USB port, or via the CMC as virtual media.
6. Select Other as the OS type, and click Next.
7. Select BIOS as the boot method and click Finish to reboot the server and start the installation.

Installing the operating system

Continue the operating system installation process by rebooting the system and configuring the installation options. Below we review specific steps to complete this task. Perform these steps on all servers.

1. On the Welcome to Oracle Linux Server 6.2 screen, select Install or upgrade an existing system, and press Next.
2. On the Oracle Linux 6 screen, click Next.
3. On install-language screen, select the installation language, and click Next.
4. On the keyboard-selection screen, select your keyboard type, and click Next.
5. On the Storage Devices screen, select Basic Storage Devices, and click Next.
6. If the Storage Device Warning pop-up appears, select Yes to use this disk.
7. On the next screen, enter the hostname (e.g., `node01`), and click Next.
8. On the Time zone screen, select your time zone from the list, check System clock uses UTC, and click Next.

9. On the root password screen, enter the administrator password (twice), and click Next.
10. On the What type of installation you would like screen, select Replace Existing Linux System(s), and click Next.
11. On the Writing storage to disk screen, format the disk and create the file systems by selecting Write changes to disk.
12. On the installation type screen, keep the default (Basic Server), and click Next to start the installation.
13. When installation completes, click Reboot.

Clustering prerequisites

Security and Firewall

Next, disable SELinux and the firewall. Perform the following steps on all servers in your cluster.

1. Open a console window and log in as root.
2. Disable SELinux by editing the file /etc/selinux/config and modifying the line

   ```
   SELINUX=enforcing
   ```

   to

   ```
   SELINUX=disable
   ```
3. Disable the firewall:

   ```
   # chckconfig iptables off
   # chkconfig ip6tables off
   ```

Modifying the kernel

After disabling the firewall, modify the kernel by performing the following steps on all servers.

1. Set the OS to boot into the non-UEK kernel by executing the command:

   ```
   # echo "savedefault --default=1 --once" | \
      grub --batch
   ```
2. Reboot the server and log in as root.
3. Remove the UEK kernel:

   ```
   # yum remove kernel-uek
   ```

Configuring networking

For our setup, we configured the two NICs on the chassis' A1 and A2 fabrics to be on independent private networks. We configured their configuration files, ifcfg-em1 and ifcfg-em2 in /etc/sysconfig/network-scripts, similar to the below (perform these steps on all servers):

```
#File1
DEVICE=em1
HWADDR=5C:F9:DD:11:0D:2F
ONBOOT=yes
IPADDR=10.152.57.102
PREFIX=24

#File2
DEVICE=em2
HWADDR=5C:F9:DD:11:0D:32
ONBOOT=yes
IPADDR=10.152.56.102
PREFIX=24
```

In our topology, we configured the next two NICs in a bond together for the public network. We used these NICs for iSCSI traffic, but do not expose those IP addresses here. We configured the files (ifcfg-em3, ifcfg-em4) as follows:

```
#File1
DEVICE=p1p1
HWADDR=5C:F9:DD:11:0C:45
ONBOOT=yes
MASTER=bond0
SLAVE=yes

#File2
DEVICE=p1p2
HWADDR=5C:F9:DD:11:0C:48
ONBOOT=yes
MASTER=bond0
SLAVE=yes

DEVICE=bond0
ONBOOT=yes
IPADDR=10.152.55.102
PREFIX=24
```

Load the NIC bond by executing the module file /etc/modprobe.d/bond.conf:

```
# echo "alias bond0 bonding" > \
   /etc/modprobe.d/bond.conf
```

Configuring DNS

For DNS resolution, add the nameserver and domain name to the file /etc/resolv.conf on each server using the following commands:

```
nameserver WWW.XXX.YYY.ZZZ
domain DOMAINNAME
```

Configuring repositories to download and install packages

Perform the following steps on all servers:

1. Configure the yum package manager to use the Oracle public software repository.
2. Download and install the Oracle Linux repo file:

```
# wget http://public-yum.oracle.com/public-yum-ol6.repo
```

3. Modify the file to use only packages from the 6.2 distribution. Edit public-yum-ol6.repo and change the line `enabled=1` to `enabled=0`. In the 6.2 stanza, change `enable=0` to `enable=1`.
4. Copy the modified file /etc/yum.repos.d:

```
# cp public-yum-ol6.repo /etc/yum.repos.d
```

5. Install the iscsi-initiator-utils package and the Grid/Oracle prerequisite packages:

```
# yum install iscsi-initiator-utils binutils gcc \
   gcc-c++ glibc compat-libcap1 \
   compat-libstdc++-33 glibc-devel libaio \
   libaio-devel libgcc libstdc++ libstdc++-devel \
   sysstat make ksh elfutils-libelf-devel \
   unixODBC unixODBC-devel

# yum install compat-libstdc++-33.i686 \
   glibc-devel.i686 libaio.i686 libaio-devel.i686 \
   libstdc++.i686 libstdc++-devel.i686 glibc.i686 \
   libgcc.i686 unixODBC.i686 unixODBC-devel.i686
```

Configuring cluster IP addresses

Enter the IP addresses for the cluster's virtual IP addresses as well as the IP addresses for each server's public and private interfaces into /etc/host. For example (Perform these steps on all servers):

```
# Public
10.152.55.101   node01
10.152.55.102   node02
…
# VIPs
10.152.55.111   node01-vip
10.152.55.112   node02-vip
…
# Private #1
10.152.56.101   node01-priv1
10.152.56.102   node02-priv1
```

```
…
# Private #2
10.152.57.101   node01-priv2
10.152.57.102   node02-priv2
…
```

Configuring NTP

Configure the NTP time service for the OS and for the Oracle Grid prerequisite on each server:

1. Configure NTP corrections to use slewing by editing the file /etc/sysconfig/ntpd and changing the line:

   ```
   OPTIONS="-u ntp:ntp  -p /var/run/ntpd.pid -g"
   ```

   to

   ```
   OPTIONS="-u ntp:ntp  -p /var/run/ntpd.pid -g -x"
   ```

2. Modify the NTP servers in /etc/ntp.conf as needed for your organization.
3. Reset the server's time:

   ```
   # service ntpdate start
   ```

4. Start the NTP daemon:

   ```
   # chkconfig ntpd on
   # service ntpd start
   ```

Configuring iSCSI with Dell EqualLogic

Perform these steps on each server.

1. Install the EqualLogic Host Integration Tools for Linux. Copy the distribution ISO to the server and mount it:

   ```
   # mount -o loop \
      equallogic-host-tools-1.1.0-1.iso /mnt
   ```

2. Run the installer and accept the defaults, including the automatic Fixes:

   ```
   # /mnt/install
   ```

3. Configure iSCSI to access volumes of the storage array. Determine secondary MAC addresses for the two interfaces, p1p1 and p2p2, that will communicate with the array. These MAC addresses may be found in the output of `iscsiadm -m iface`. Often they are simply related to their primary MAC address (found in the output of ifconfig, e.g.); namely, the last (least significant) hexadecimal pair in the secondary MAC addresses is one more than the corresponding pair in the primary MAC address. You can check the server's iDRAC to determine each interface's primary and secondary MAC address. Assuming the secondary MAC addresses for the

two NICs are bnx2i.XX:XX:XX:XX:XX:YY and bnx2i.XX:XX:XX:XX:XX:ZZ, assign IP addresses to these devices on the storage array's subnet

```
# iscsiadm -m iface -I bnx2i.XX:XX:XX:XX:XX:YY \
   -o update -n iface.ipaddress  -v 192.168.5.XX1

# iscsiadm -m iface -I bnx2i.XX:XX:XX:XX:XX:ZZ \
   -o update -n iface.ipaddress -v 192.168.5.YY1
```

4. Modify the EqualLogic HIT configuration file , /etc/equallogic/eql.conf, to change the iSCSI transport from TCP to that used by the NICs. Edit the file and change the line:

   ```
   IscsiInitiator = tcp
   ```

   to

   ```
   IscsiInitiator = bnx2i
   ```

5. Restart EqualLogic HIT daemon:

   ```
   # service ehcmd restart
   ```

6. Probe the array's iSCSI portal to find targets. Here, the portal address is 192.168.5.11. Specify the iSCSI names of both interfaces:

   ```
   # iscsiadm -m discovery -t st -p 192.168.5.11 \
      -I bnx2i.XX:XX:XX:XX:XX:YY \
      -I bnx2i.XX:XX:XX:XX:XX:ZZ
   ```

7. List the targets:

   ```
   # iscsiadm -m node
   ```

8. Login into each target (e.g., iqn.2001-05.com.equallogic:4-52aed6-df2c63f66-6ef0000001e500b6-l3):

   ```
   # ehcmcli login -T \
      iqn.2001-05.com.equallogic:4-52aed6-df2c63f66-
   6ef0000001e500b6-l3
   ```

   Note: The previous command prints the name of the iSCSI volume's block device (a symbolic link in the directory /dev/eql).

9. On the first server only, create one partition on each device:

   ```
   # fdisk -cu /dev/eql/NAME_OF_VOLUME
   ```

   a. Enter `n`
   b. Enter `p`
   c. Enter `1`
   d. Press Enter to select the default starting sector.
   e. Press Enter to select the default ending sector.
   f. Enter `w`

Configuring Oracle installation prerequisites

1. Create the following groups and logins for Oracle Grid and Database. Replace the numeric user IDs and group IDs with the corresponding used by your organization:

```
# groupadd -g 1001 oinstall
# groupadd -g 1002 dba
# groupadd -g 1003 asmadmin
# groupadd -g 1004 asmdba
# useradd -u 1003 -g oinstall \
   -G dba,asmadmin,asmdba grid
# useradd -u 1002 -g oinstall \
   -G dba,asmadmin,asmdba oracle
```

2. Assign passwords to the oracle and grid accounts:

```
# passwd oracle
# passwd grid
```

3. Set the umask for the oracle and grid accounts:

```
# echo "umask 022" >> ~oracle/.bash_profile
# echo "umask 022" >> ~grid/.bash_profile
```

4. Create the directory for Oracle software, inventory and other data:

```
# mkdir –p /u01/app/oracle
# mkdir /u01/app/grid
# chown –R oracle:oinstall /u01
# chmod –R g+w /u01/app
```

5. Increase the resource limits for the oracle and grid accounts by editing the file /etc/security/limits.conf and adding the lines to the end of the file:

6.
```
oracle - nproc 16384
oracle - nofile 65536
oracle - stack 10240
grid - nproc 2047
grid - nofile 65536
```

   Modify the system configuration per Oracle prerequisite by adding the following lines to /etc/sysctl.conf:

```
kernel.msgmnb = 65536
kernel.msgmax = 65536
kernel.sem = 250 32000 100 128
fs.file-max = 6815744
fs.aio-max-nr = 1048576
net.ipv4.ip_local_port_range = 9000 65500
net.core.rmem_default = 262144
net.core.wmem_default = 262144
```

```
net.core.rmem_max = 4194304
net.core.wmem_max = 1048576
```

7. Compute the values of the SYSV shared-memory parameters, which depend upon the amount of server RAM.
   a. The value of the kernel.shmall parameter is the size of system RAM divided by the page size. The page size is usually 4096 bytes and can be determined by the command:
      ```
      # getconf PAGE_SIZE
      ```
   b. For RAM equal to 128 GB, the kernel.shmall parameter equals 128*1024*1024*1024/4096 or 33,554,432. If the RAM is 32 GB, the kernel.shmall is proportionally smaller or 8,388,608. Set kernel.shmall to the correct value for your system by adding a line similar to the following to the end of /etc/sysctl.conf:
      ```
      kernel.shmall = 33554432
      ```
   c. The value of the kernel.shmmax parameter is one-half the size of total memory. So for our 128 GB example, this parameter should be set to 128*1024*1024*1024/2 or 68,719,476,736; for 32 GB, kernel.shmmax, is proportionally smaller: 32*1024*1024*1204/2 = 17,179,869,184. Add a line similar to the following to the end of /etc/sysctl.conf:
      ```
      kernel.shmmax = 68719476736
      ```
8. To effect these changes without a reboot, run the command:
   ```
   # sysctl -p
   ```
9. Create the file /etc/init/ohasd.conf, which starts the Oracle Grid HA daemon:
   ```
   #start init.ohasd
   start on runlevel [35]
   stop on runlevel [016]
   respawn
   exec /etc/init.d/init.ohasd run >/dev/null 2>&1
   </dev/null
   ```
10. Repeat steps 1-9 for each server. Note that several of the configuration files can be copied from server to server.

INSTALLING AND CONFIGURING ORACLE RAC

Creating the cluster

Installing Oracle Grid Infrastructure

1. All nodes in the server should be configured as in the previous section and turned on.
2. Copy the Oracle Grid Infrastructure 11*g* R2 software to one of the nodes.
3. Log in as the grid account and unzip the software:

```
# su - grid
% unzip /tmp/linux.x64_11gR2_grid.zip
```

4. Run the installer after setting the X Window DISPLAY variable:

```
% grid/runInstaller
```

5. If the installer warns that an X Window utility cannot be found, type `y` to override.
6. On the Select Installation Option screen, select Install and Configure Grid Infrastructure for a Cluster, and click Next.

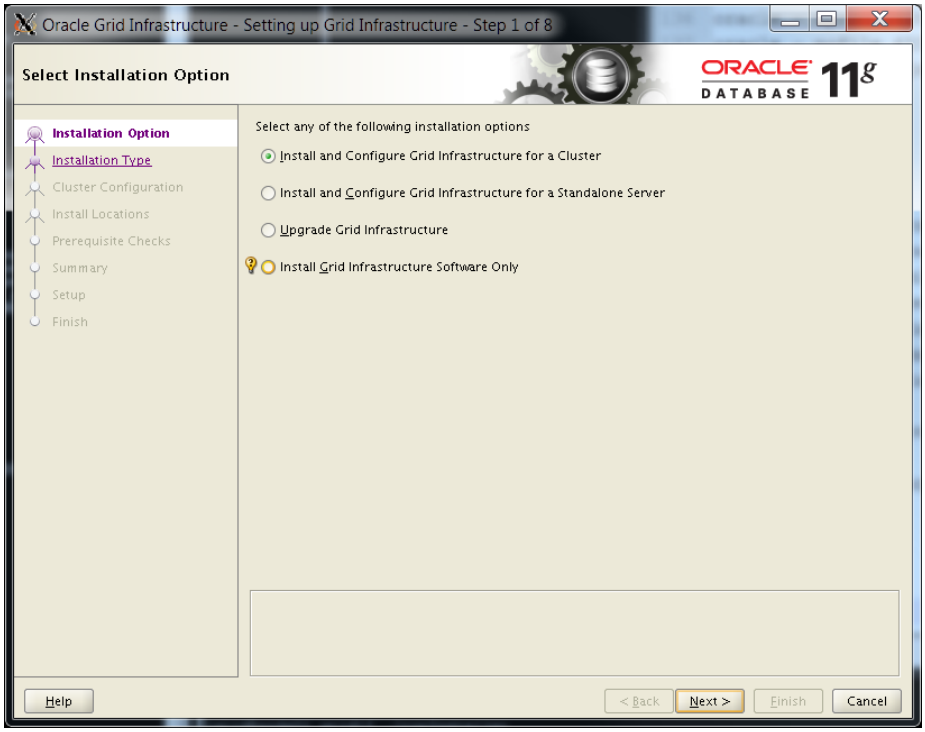

7. On the Select Installation Type screen, select Typical Installation, and click Next.

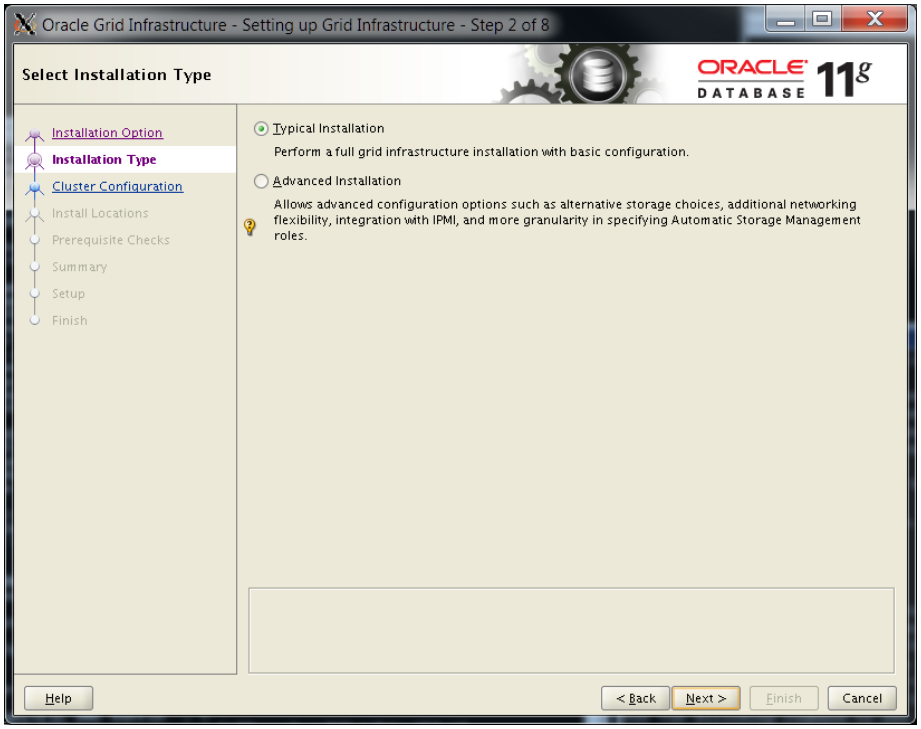

8. On the Specify Cluster Configuration screen, enter the SCAN name (cluster name in DNS).
9. On the same screen, click Add to enter the hostname and VIP name for each additional node (e.g.,Hostname: node04; Virtual IP Name: node04-vip).

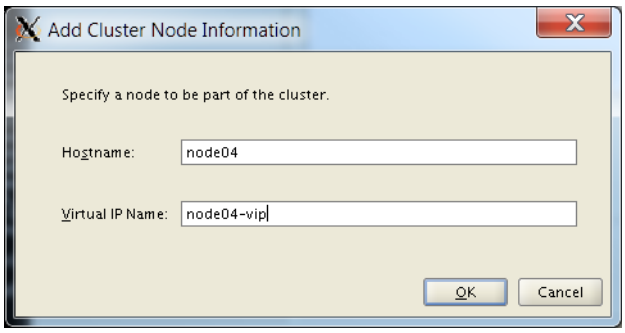

10. On the same screen, click SSH connectivity, enter the grid OS password, and on the sub-screen, click Setup.

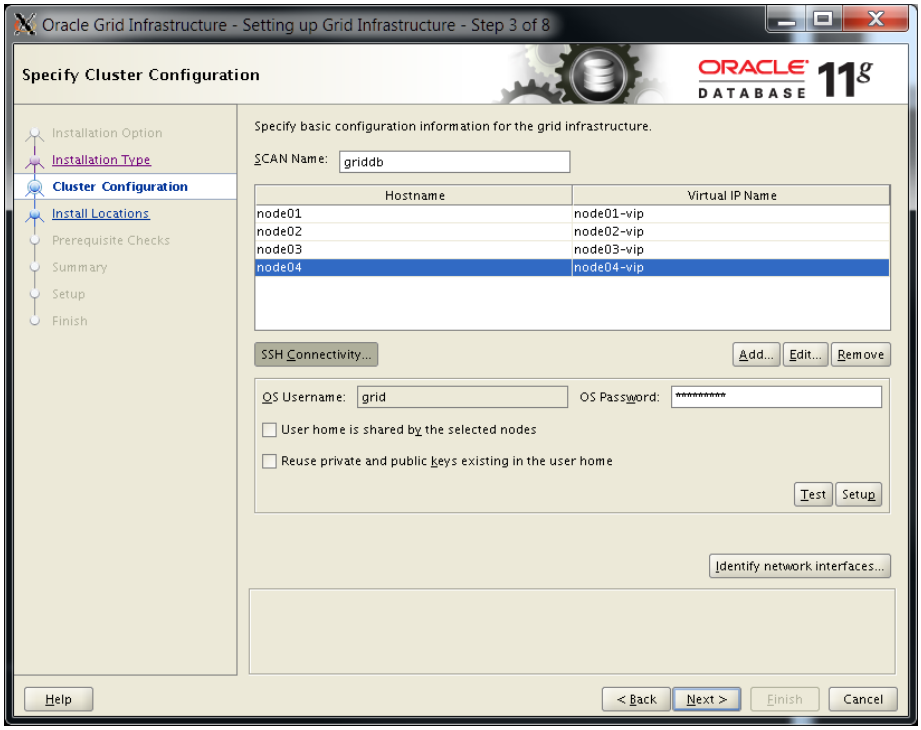

11. Click Next.

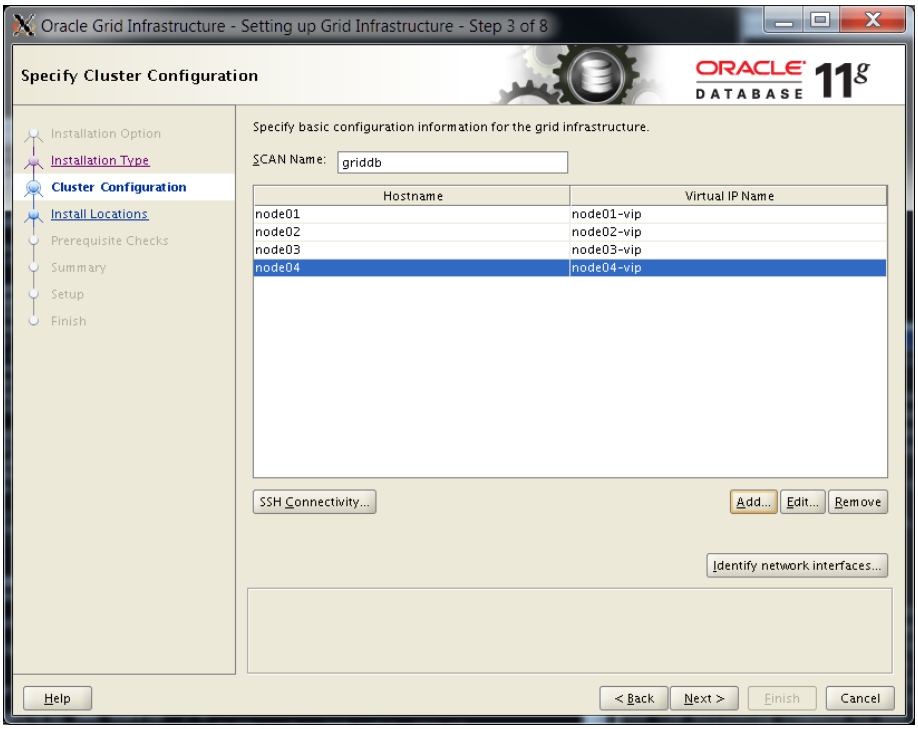

12. On the Specify Install Locations screen, enter /u01/app/grid for Oracle Base, `/u01/app/11.2.0/`grid for Software Location, select Automatic Storage Management for Cluster Registry Storage Type, enter the SYSADM password, and select asmadmin for OSASM group. Click Next.

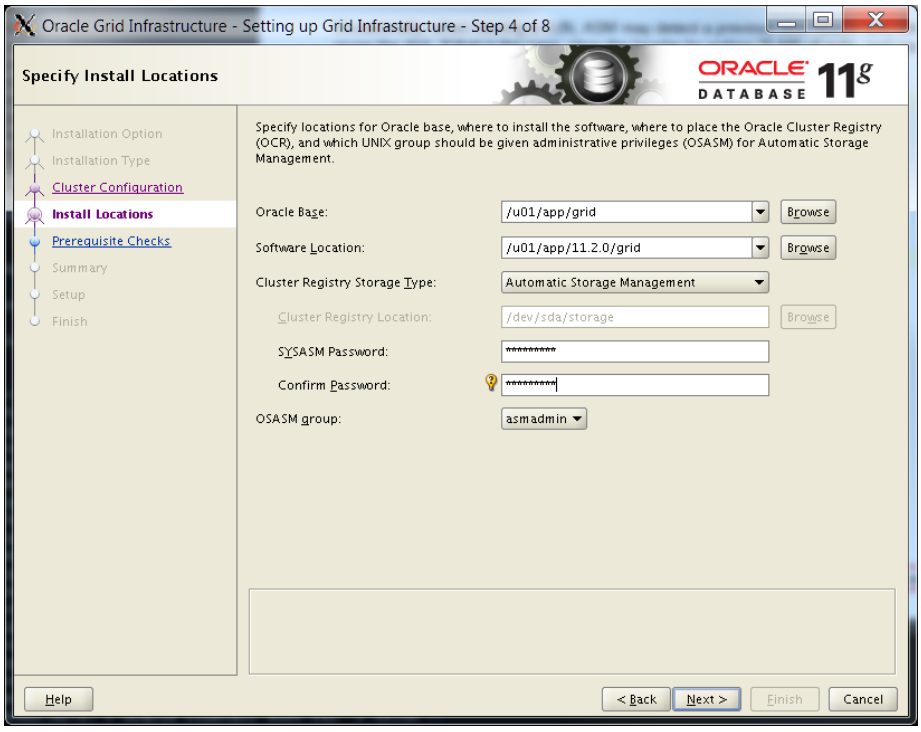

13. On the Create ASM Disk Group screen, click Change Discovery Path. Enter `/dev/eql` and click OK.

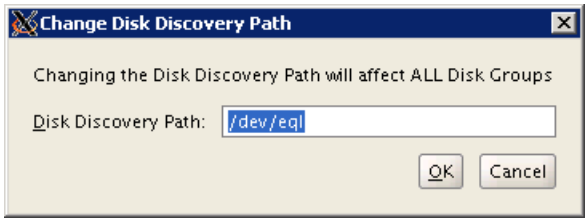

14. Three disks should populate the Disk Window. Select them, and click Next.

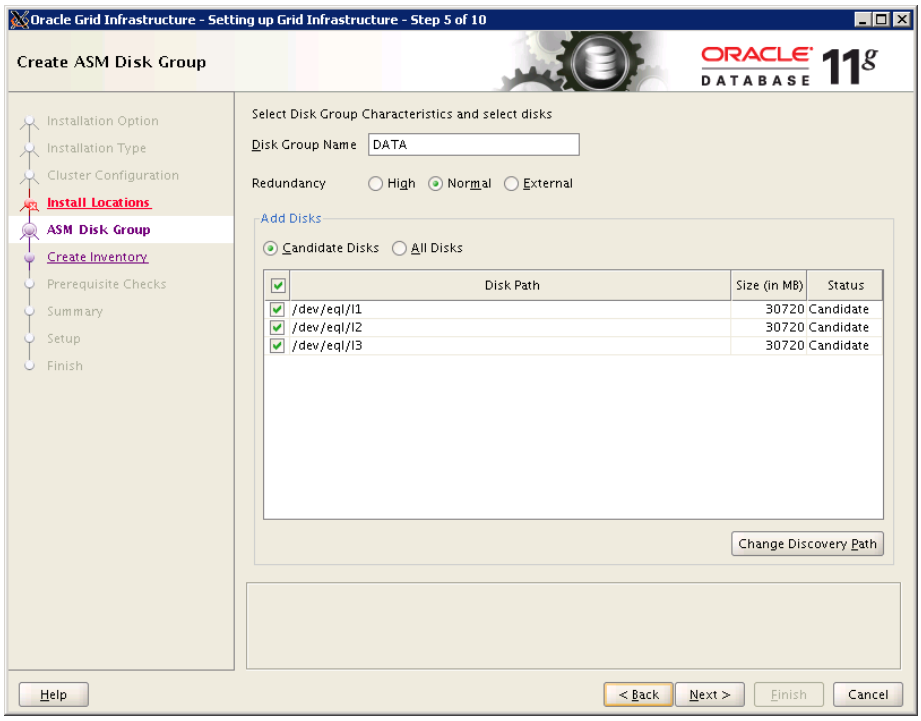

15. On the Create Inventory screen, enter `/u01/app/oraInventory` and click Next.

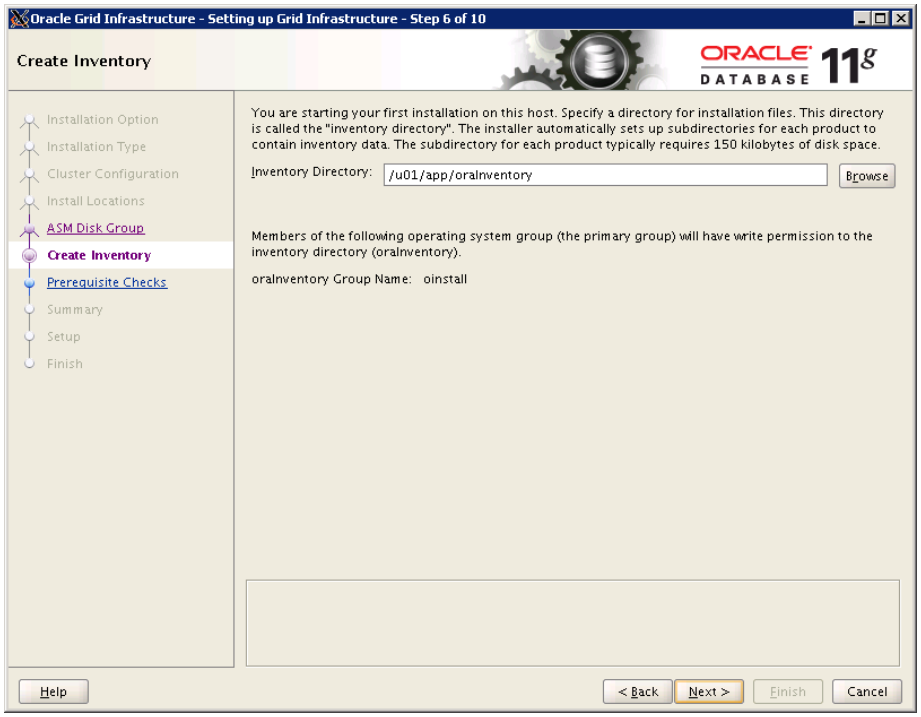

16. On the Perform Prerequisite Checks screen, examine failed checks. If any item can be fixed automatically, click Fix and Check Again. If all failed checks are missing packages and you installed the packages given above, select Ignore All, and click Next.

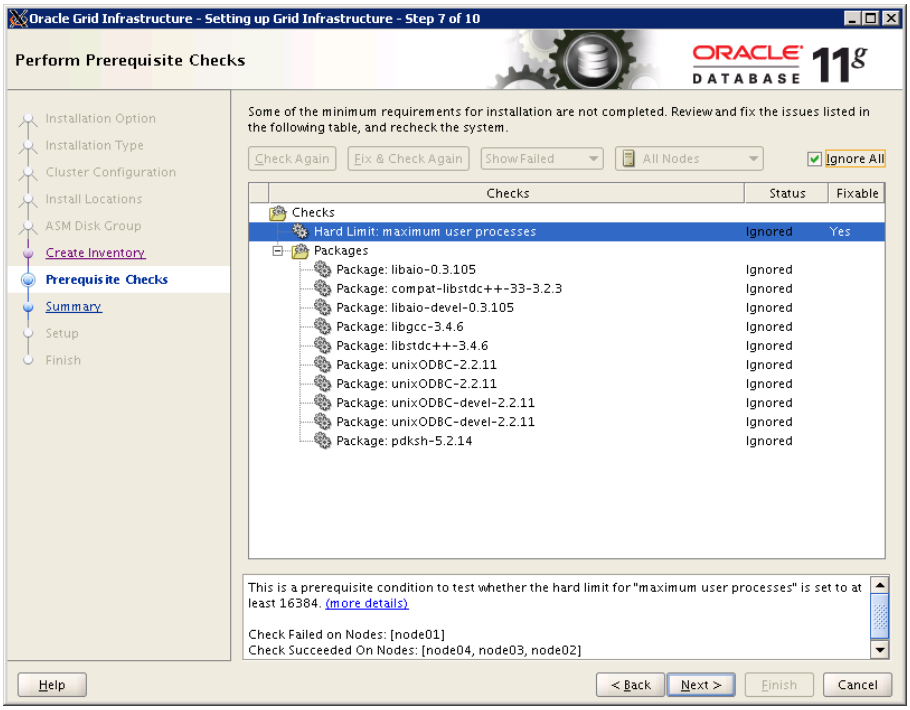

17. On the Summary screen, click Finish to begin installation.

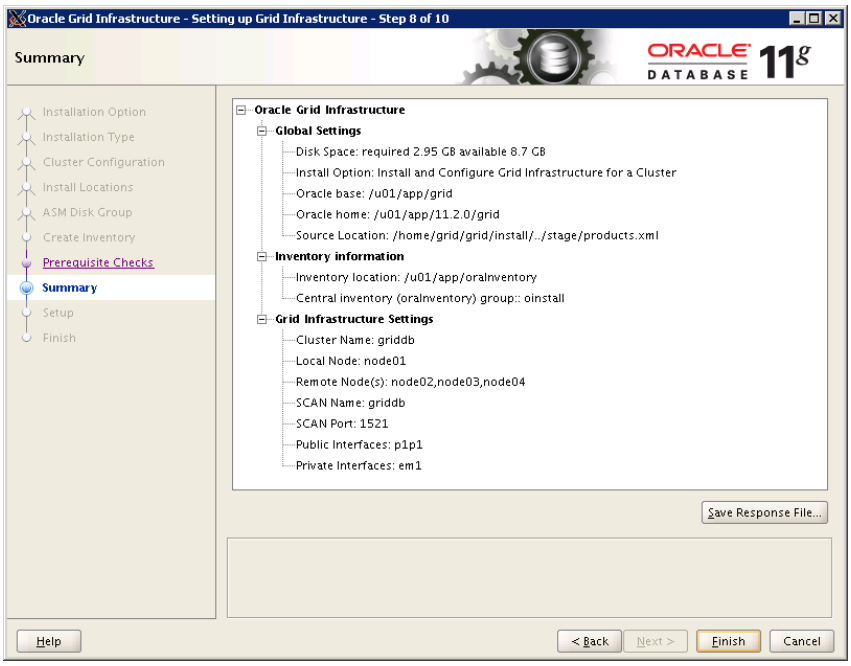

18. Toward the end of the installation, an Execute Configuration scripts window pops up. Open a second console window as root, and run the first script on each node.

```
## For node01
# /u01/app/oraInventory/orainstRoot.sh
## For each remote node
# ssh node02 /u01/app/oraInventory/orainstRoot.sh
…
```

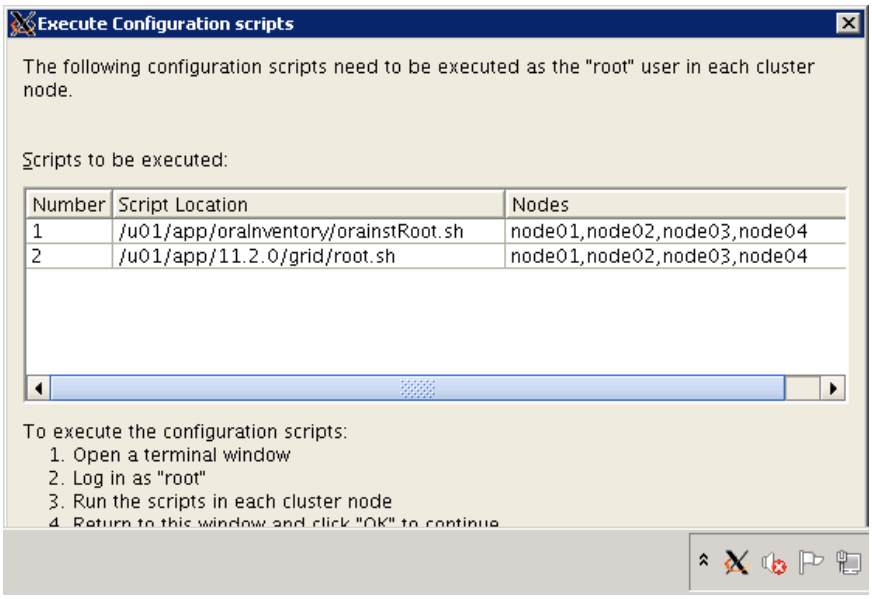

19. Run the second script in one window. When the script's output indicates that it is trying to start the ohasd daemon (which it cannot), type `initctl start ohasd` in the first console window. For example,

```
## In window 2
# ssh node02 /u01/app/11.2.0/grid/root.sh
## In window 1
# ssh node02 initctl start ohasd
```

20. After the scripts have run, click OK on the Execute Configuration scripts screen, and click Close.

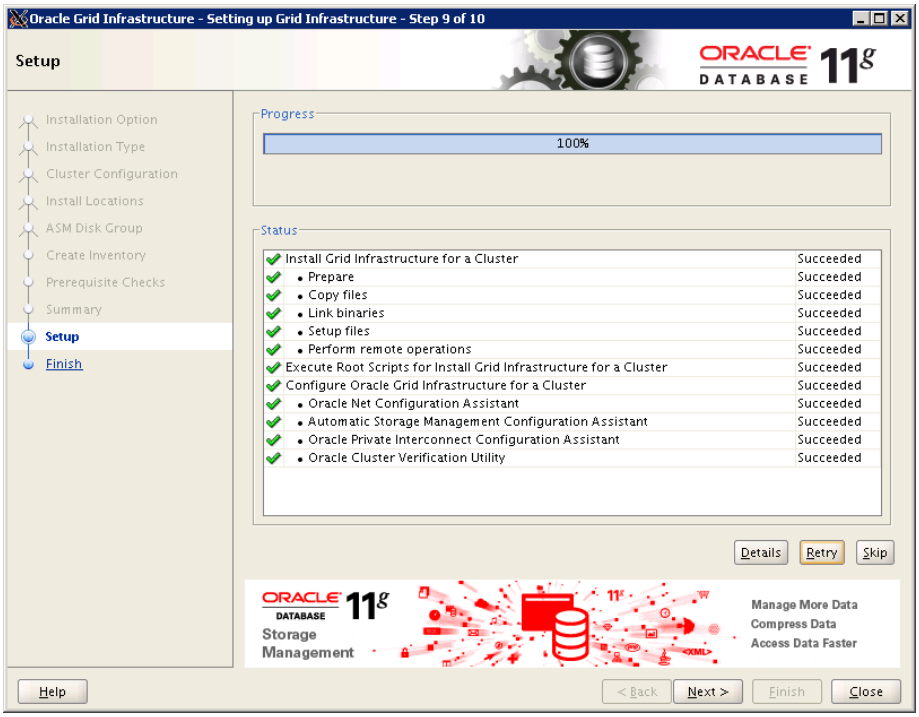

21. On the Finish screen, click Close.

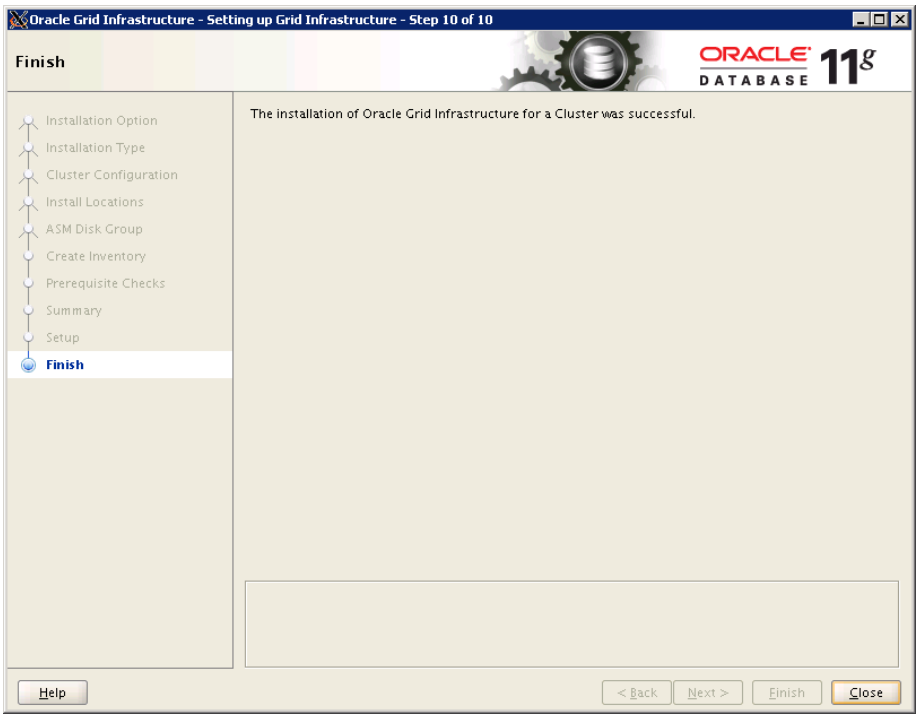

Installing Oracle RAC 11*g* (Oracle Database) on the cluster

In this section, we provide step-by-step instructions for installing Oracle RAC on the cluster.

1. All nodes in the cluster should be on.
2. Copy the Oracle Database 11*g* R2 software to one of the nodes.
3. Log in as the oracle account and unzip the software:

```
# su – oracle
% unzip /tmp/ linux.x64_11gR2_database_1of2.zip
% unzip /tmp/ linux.x64_11gR2_database_2of2.zip
```

4. Run the installer after setting the X Window DISPLAY variable:

```
% database/runInstaller
```

5. If the installer warns that an X Window utility cannot be found, type `y` to override.

6. On the Configure Security Updates screen, enter the appropriate contact information, and click Next.

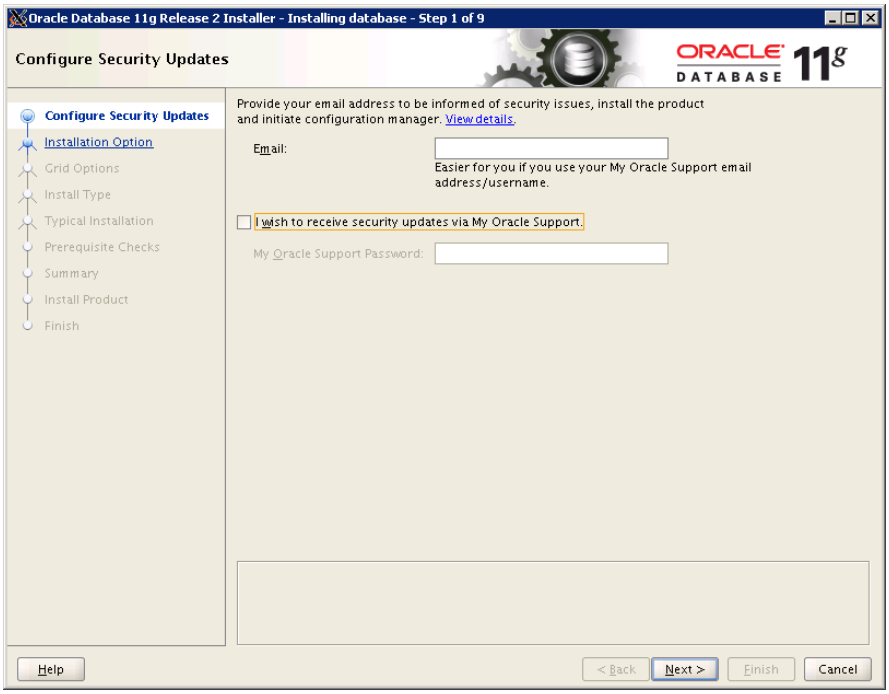

7. On the Select Installation Option screen, select Create and configure a database, and click Next.

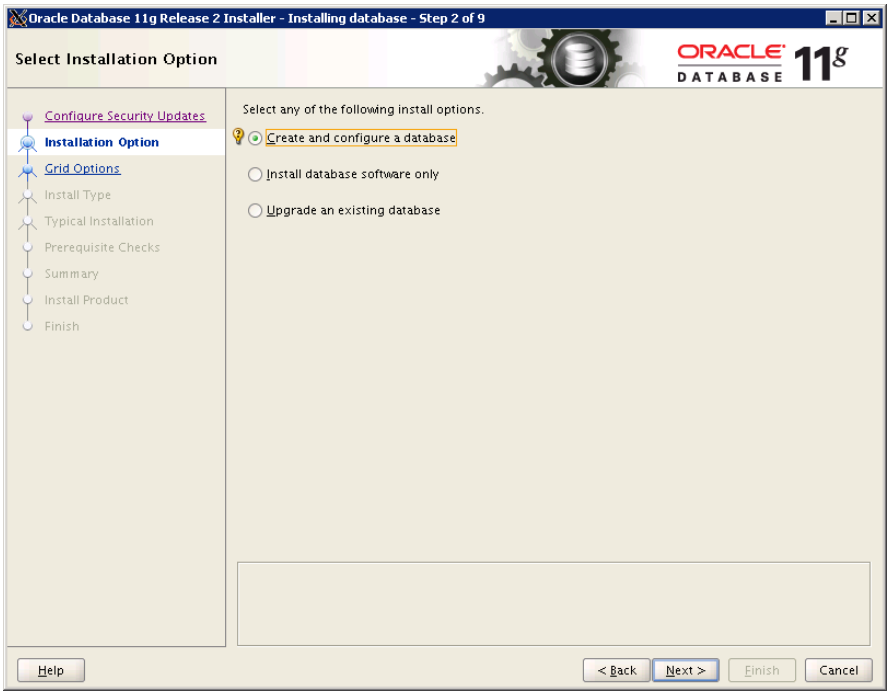

8. On the System Class screen, select Server Class, and click Next.

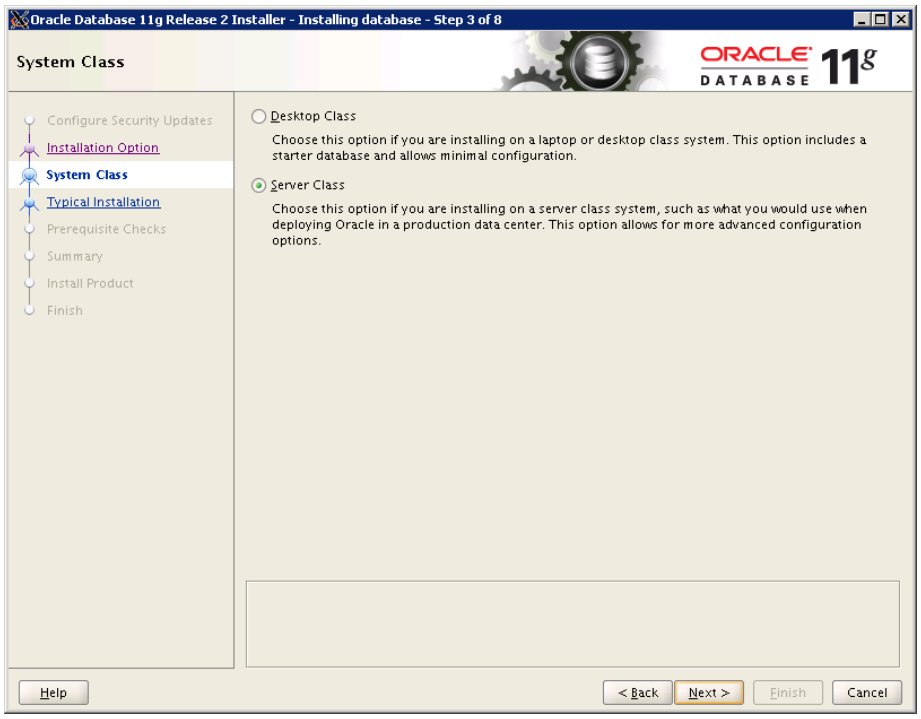

9. On the Node Selection screen, select Real Application Cluster database installation, and select all the nodes in the cluster.

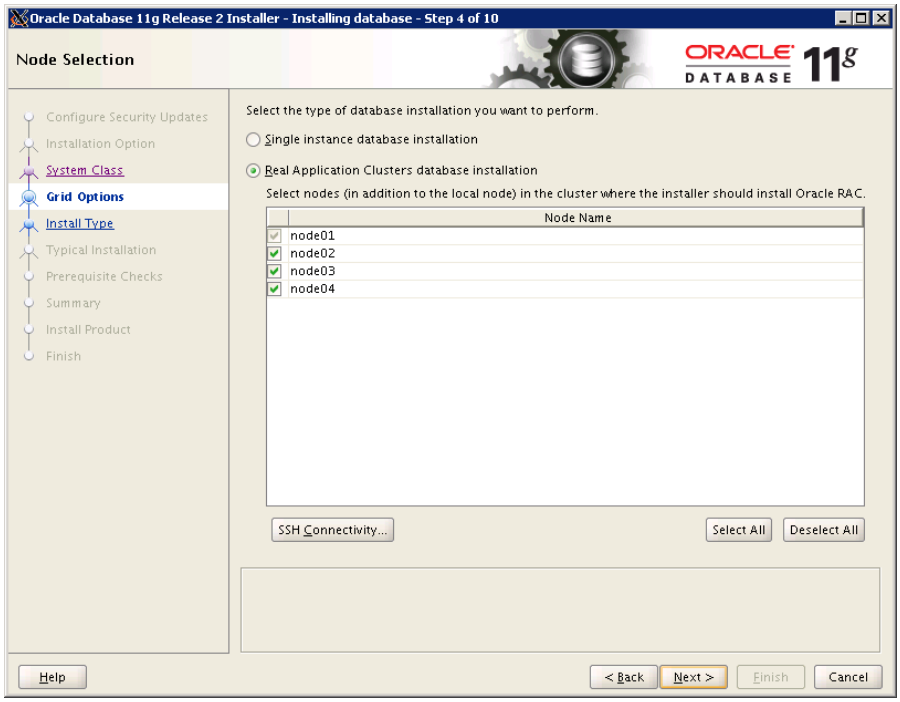

10. On the same screen, press SSH connectivity, enter the oracle OS password, and on the sub-screen, click Setup. Click Next after the configuration completes.

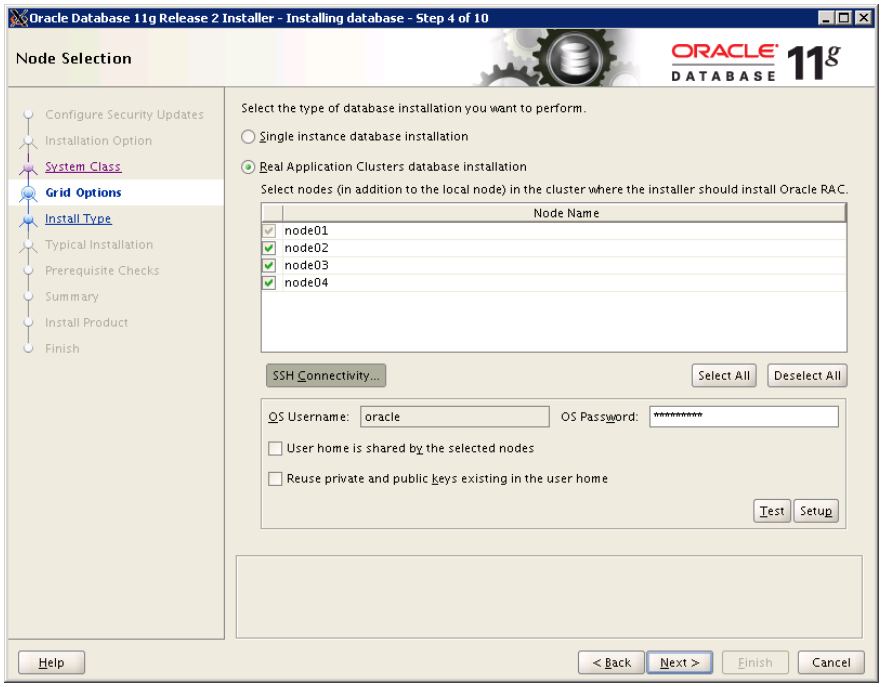

11. On the Select Install Type screen, select Typical Install, and click Next.

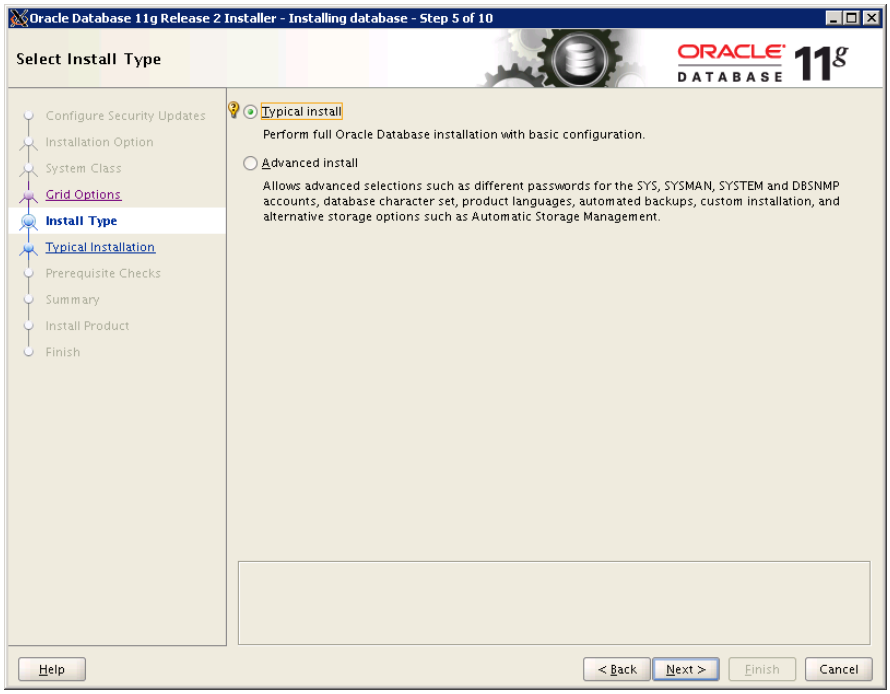

12. On the Typical Install Configuration screen, enter `/u01/app/grid` for Oracle base, enter `/u01/app/grid/product/11.2.0/db_home_1` for Software location, select Automatic Storage Management for Storage Type, enter OCR for Database file location, enter the ASMSNMP password, select Enterprise Edition for Database edition, select dba for OSDBA Group, enter the database administrative password, and click Next.

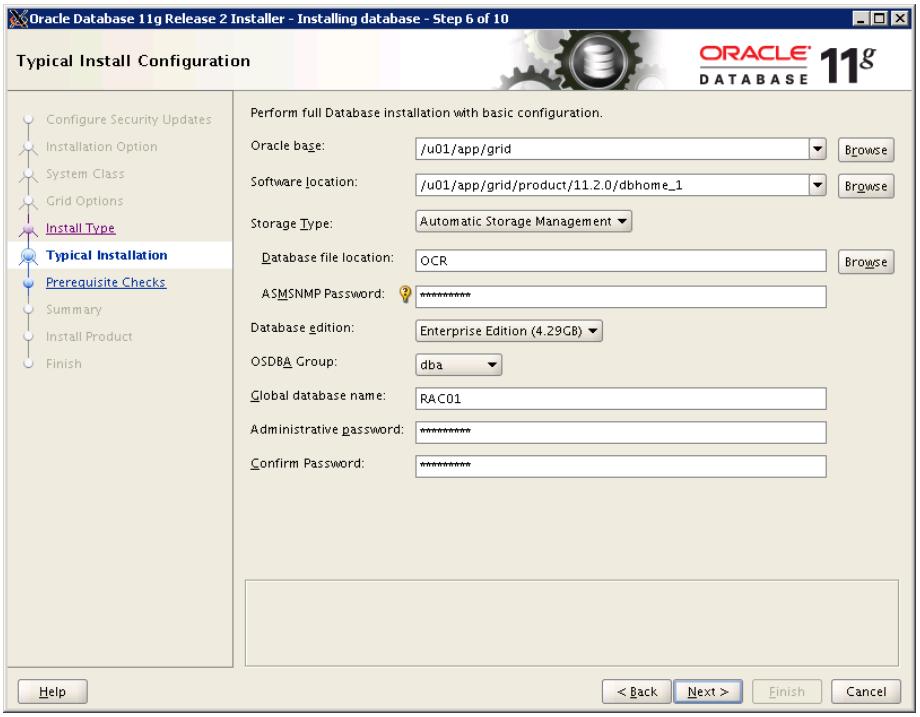

13. On the Perform Prerequisite Checks screen, select Ignore All if the only failures are missing packages and you installed all software packages above. Click Next.

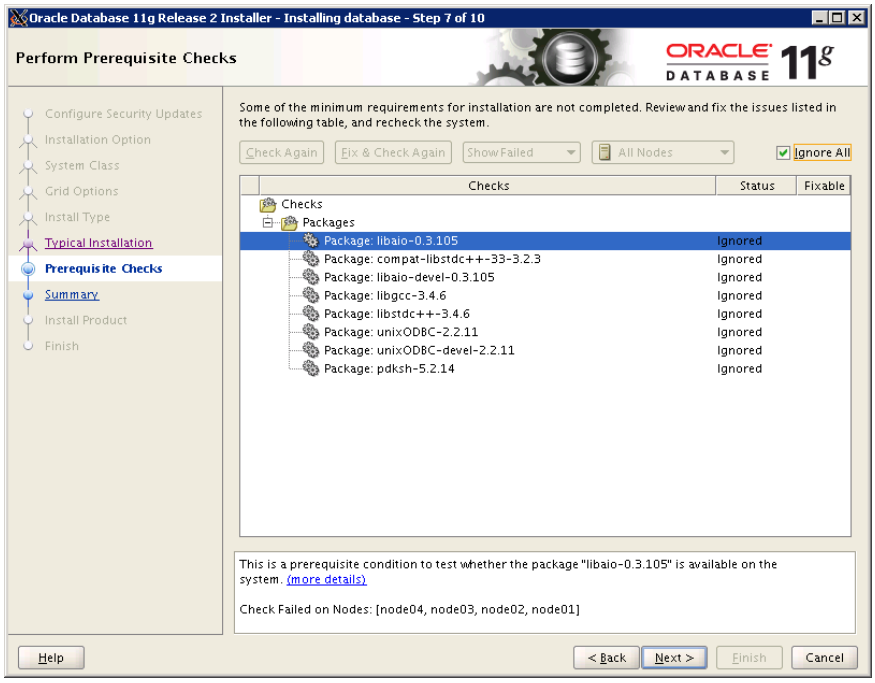

14. On the Summary screen, click Next to start the database installation.

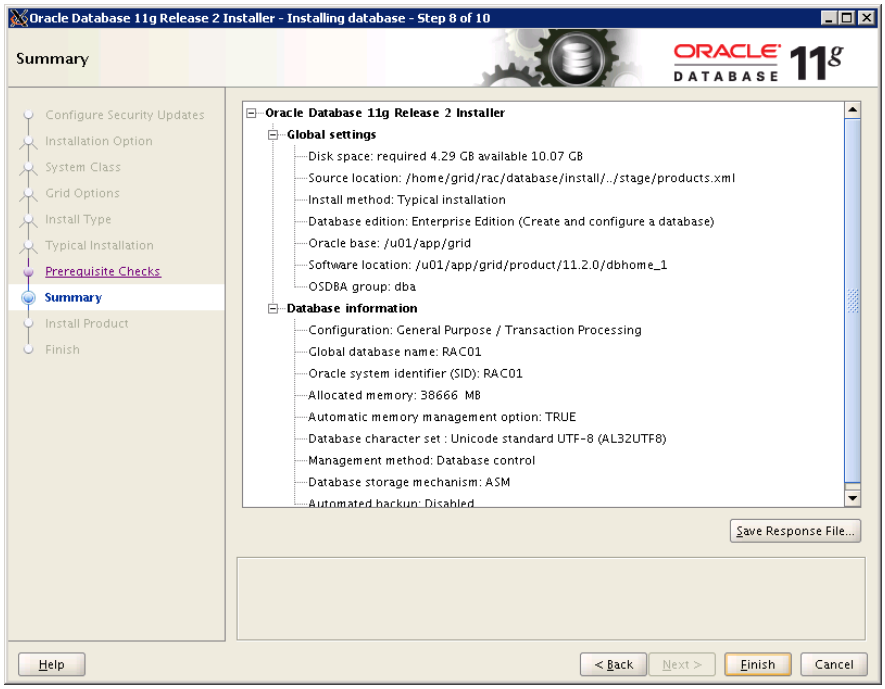

15. When the software installation completes, a database is created.

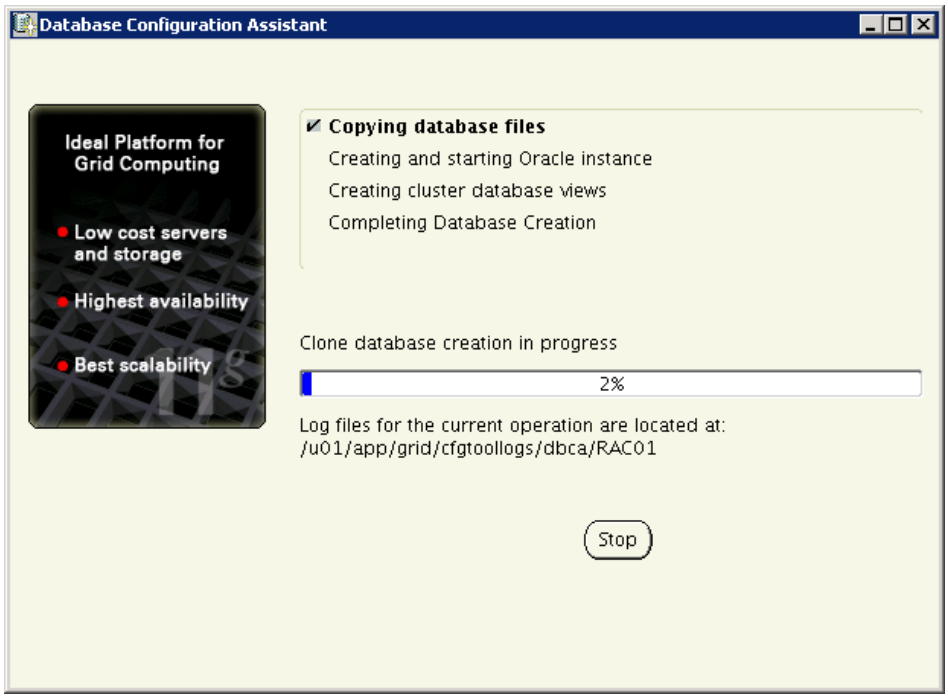

16. When the Database Configuration Assistant Warning screen appears, run the network configuration wizard, /u01/app/11.2.0/grid/bin/netca.

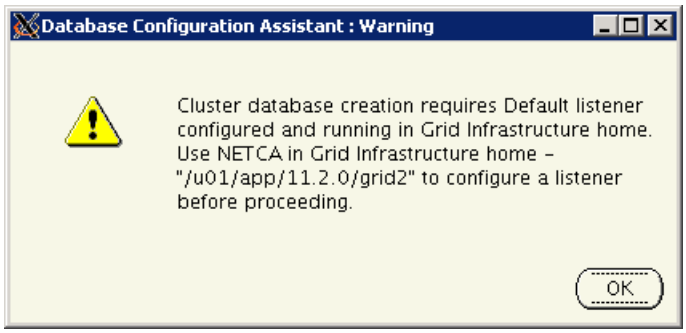

17. On the Oracle Net Configuration Assistant: Welcome screen, select Listener configuration, and click Next.

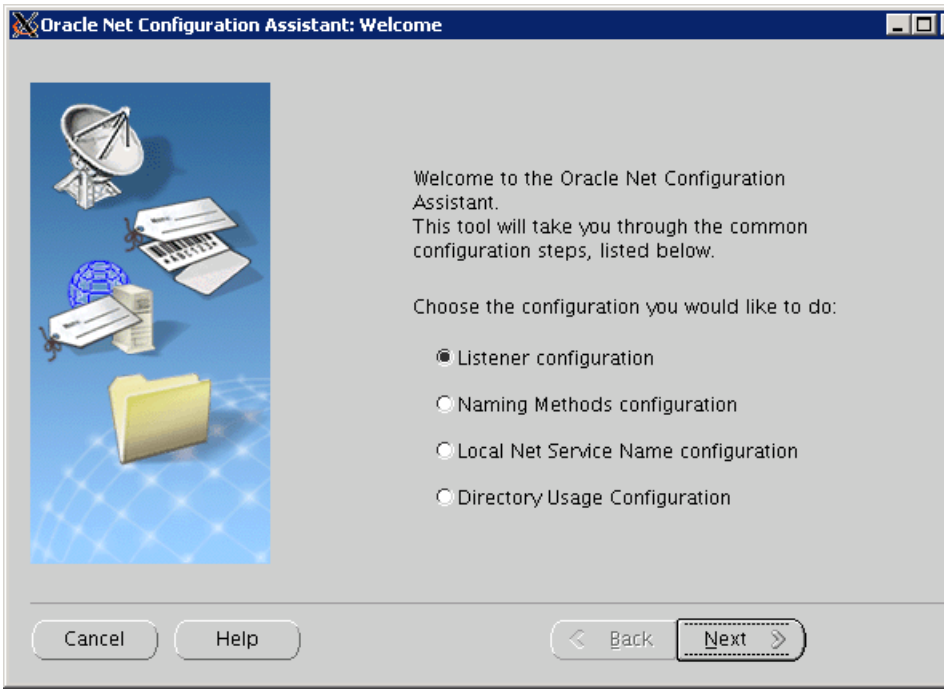

18. On the Listener Configuration, Listener screen, select Add, and click Next.

19. On the Listener Configuration, Listener Name screen, enter the Listener name, and click Next.

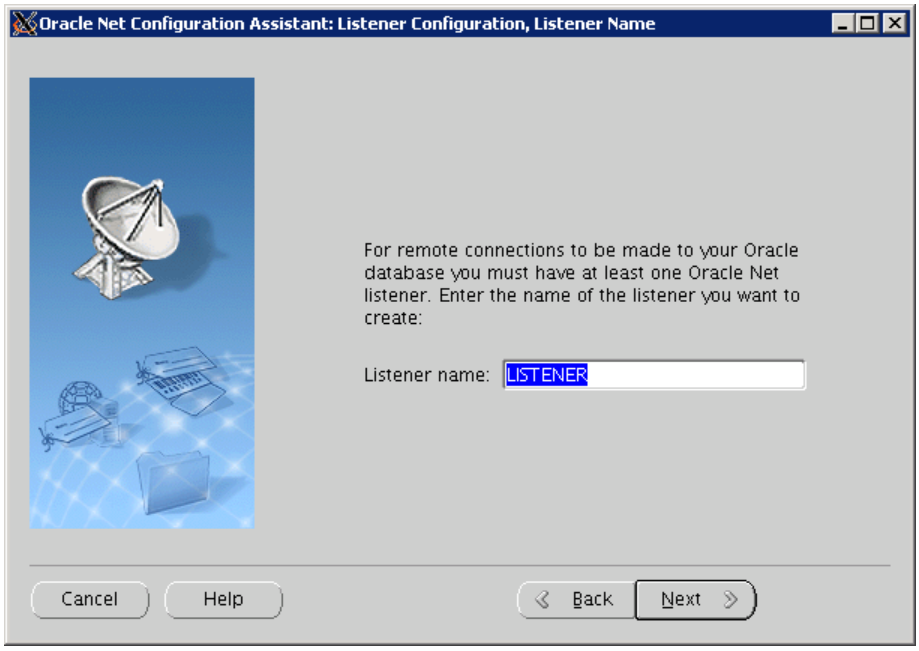

20. On the Listener Configuration, Select Protocols screen, select TCP, and click Next.

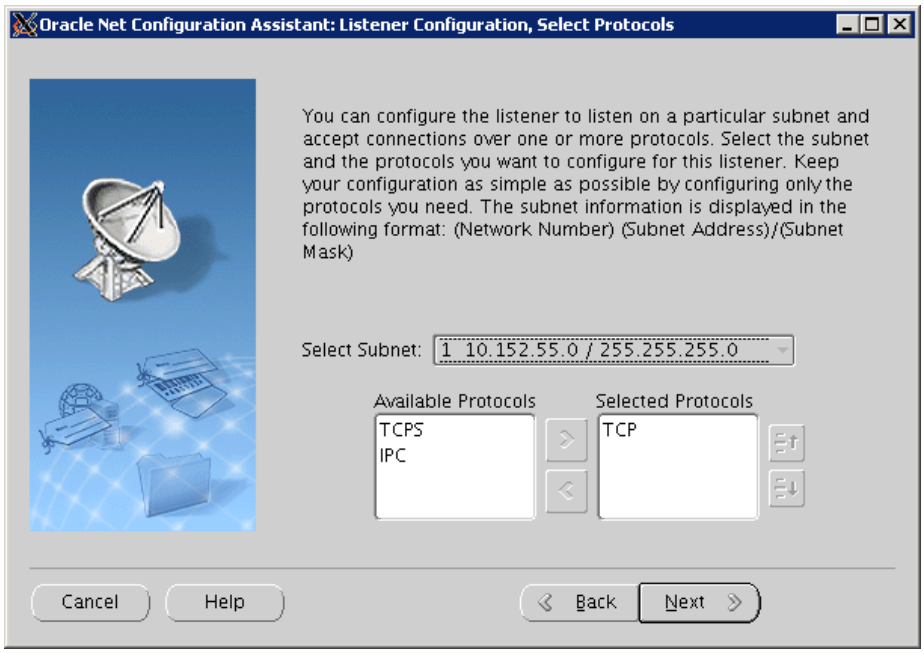

21. On the Listener Configuration, TCP/IP Protocol screen, select Use the standard port number of 1512, and click Next.

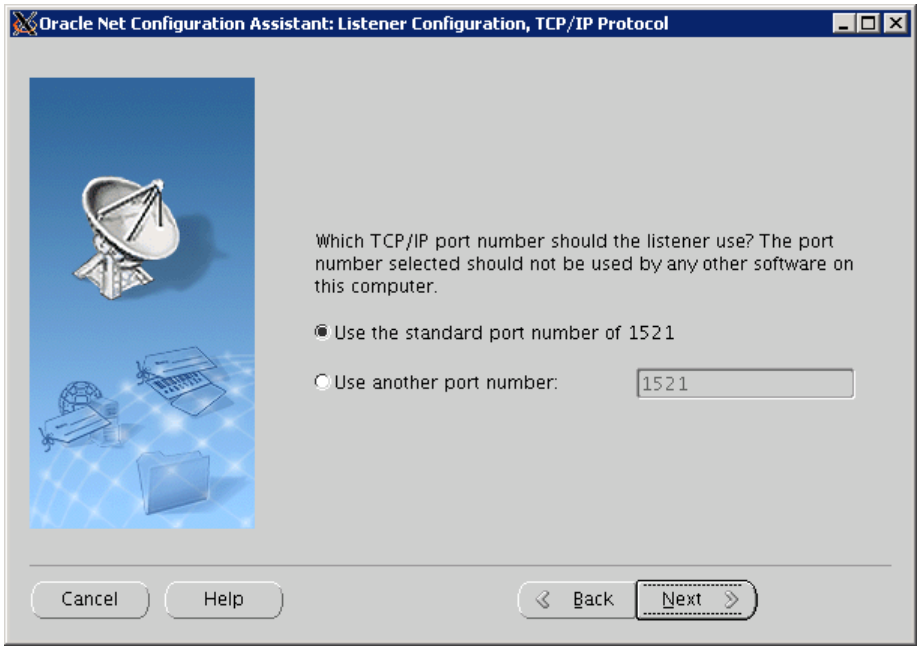

22. On the Listener Configuration, More Listeners screen, select No, and click Next.

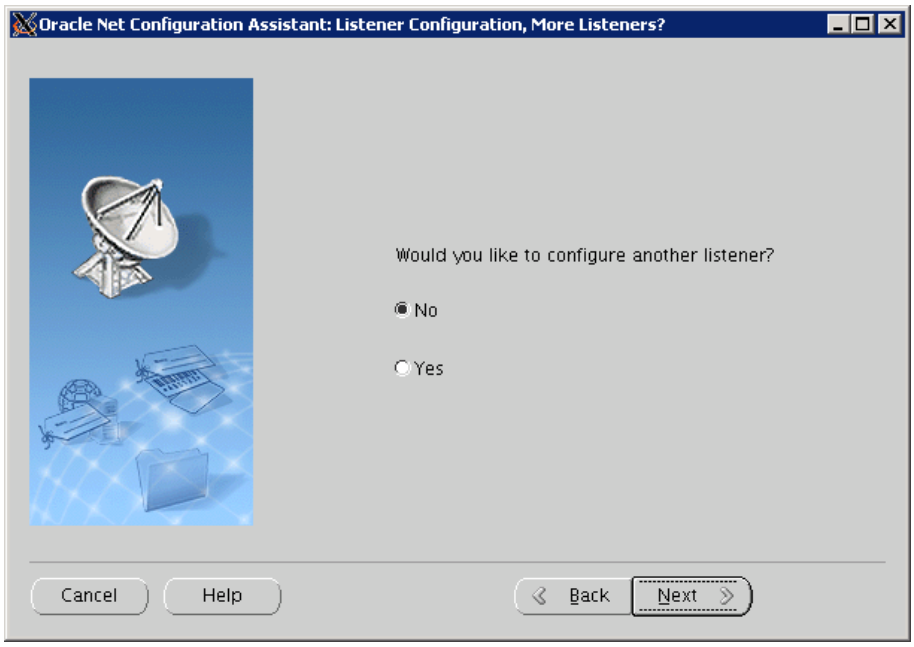

23. On the Listener Configuration Done screen, click Next.

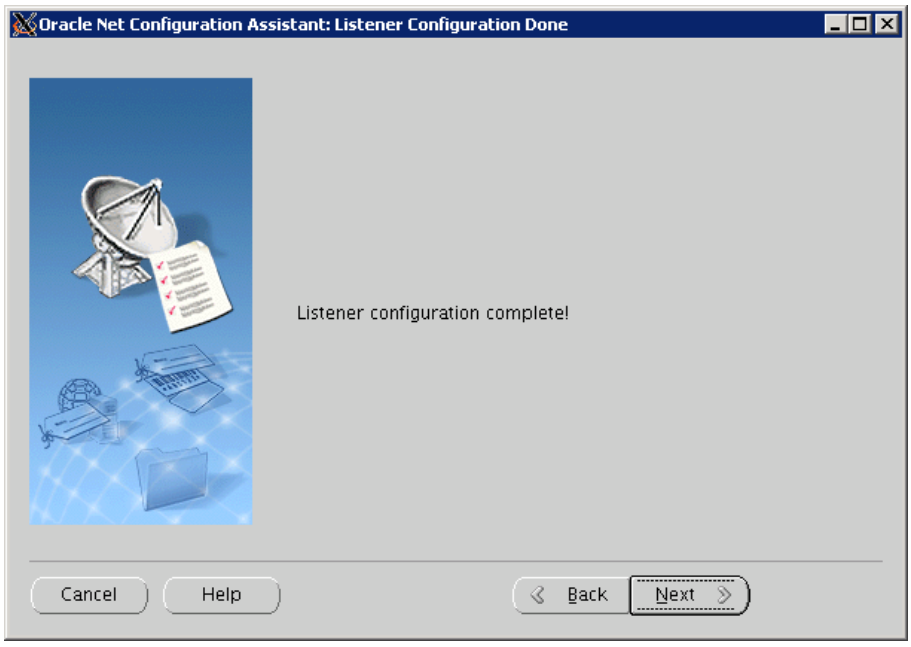

24. On the Welcome screen, click Finish.

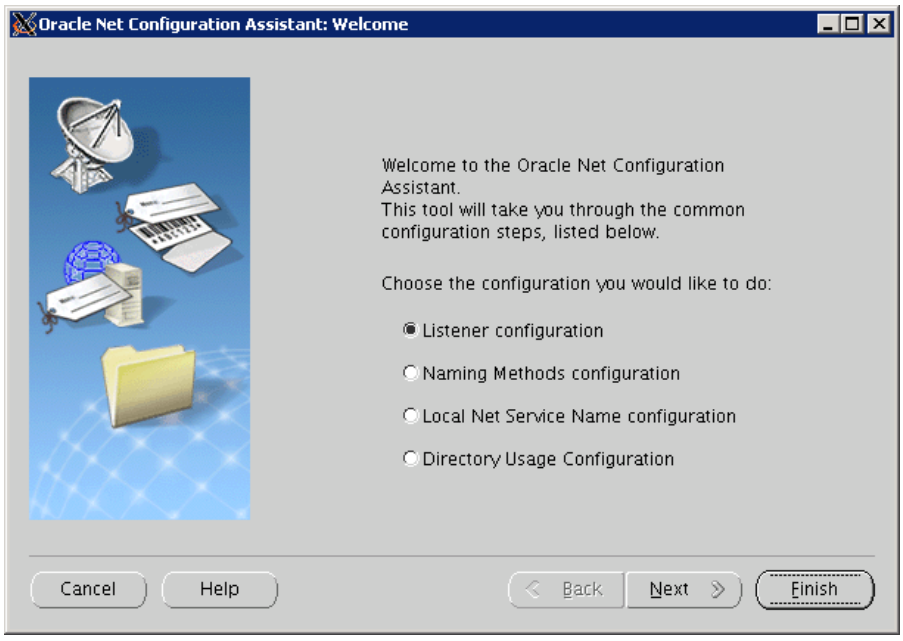

25. When database creation completes, click OK.

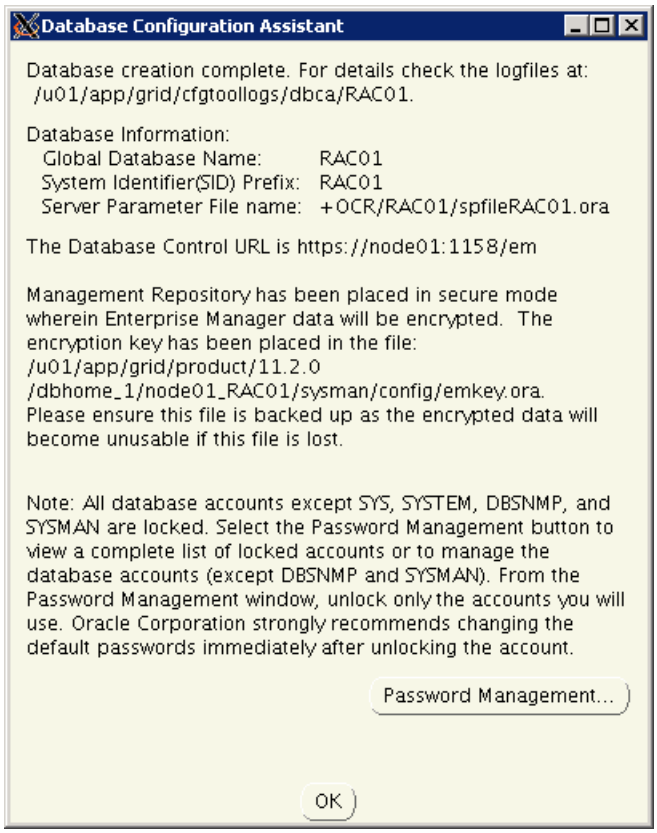

26. When the Execute Configuration scripts screen appear, open a second console window and log in as root. Run the first script on the remote nodes, and the second script on all nodes.

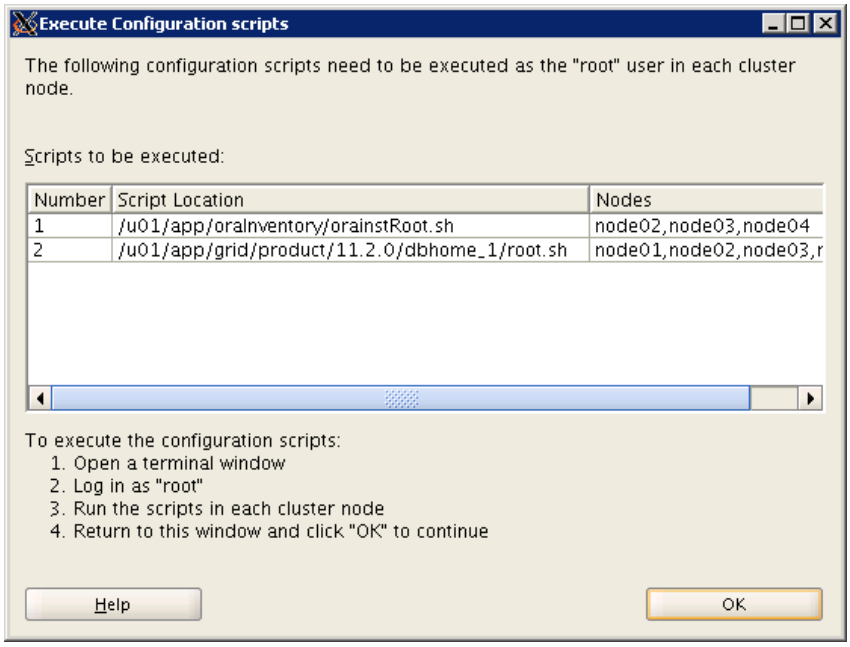


```
# ssh node02 /u01/app/oraInventory/oraInstRoot.sh
# ssh node03 /u01/app/oraInventory/oraInstRoot.sh
…
# /u01/app/grid/product/11.2.0/dbhome_1/root.sh
# ssh node02 \
   /u01/app/grid/product/11.2.0/dbhome_1/root.sh
…
```

27. After the scripts have run, click OK on the Execute Configuration scripts screen, and click Close on the Install Product screen.

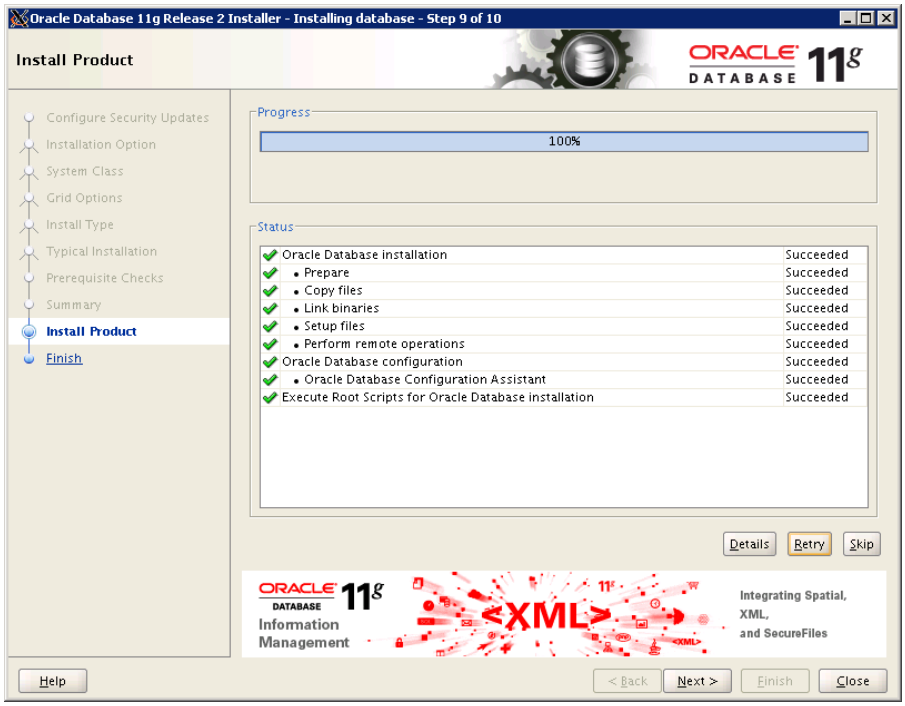

28. On the Finish screen, click Close.

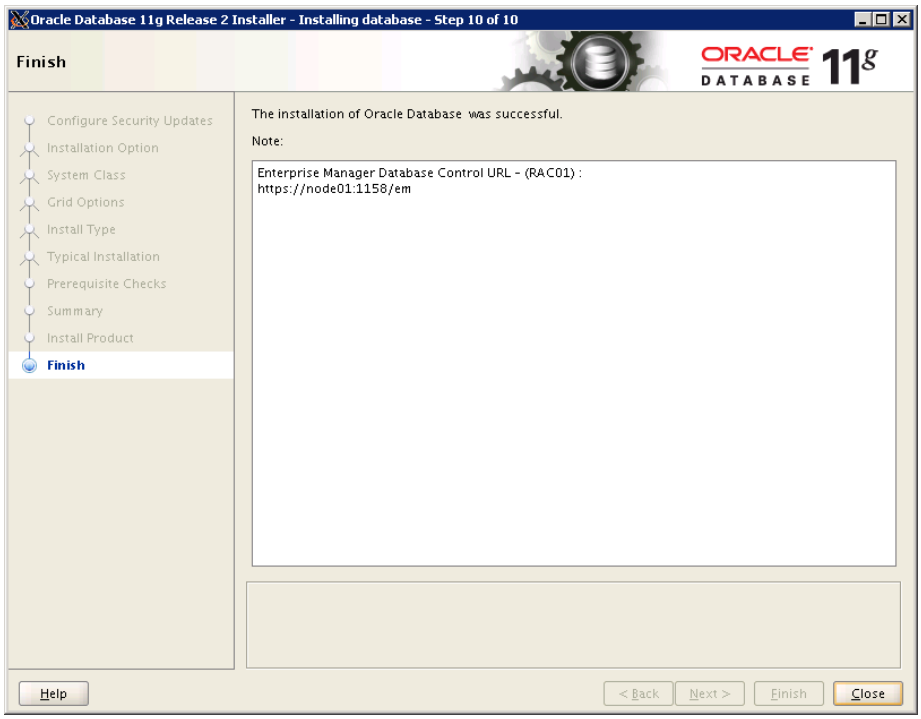

SUMMARY

The compute density of your data center is critical to your bottom line – packing more compute power into a smaller space has the potential to greatly reduce data center costs. The ultra-dense, quarter-height Dell PowerEdge M420 blade server combines with Dell PowerConnect and Force10 switches, Dell PowerEdge M1000e blade enclosure, and Dell EqualLogic PS6110XS arrays to maximize the compute density of your data center.

APPENDIX A – DELL POWEREDGE M420 SYSTEM DETAILS

Figure 5 presents the server configuration we used for this guide.

System	Dell PowerEdge M420
Power supplies	Dell PowerEdge M1000e Blade Enclosure
Total number	6
Vendor and model number	Dell A236P-00
Wattage of each (W)	2,360
Cooling fans	Dell PowerEdge M1000e Blade Enclosure
Total number	9
Vendor and model number	Dell YK776 Rev. X50
Dimensions (h x w) of each	3.1" x 3.5"
Volts	12
Amps	7
General	
Number of processor packages	2
Number of cores per processor	8
Number of hardware threads per core	2
System power management policy	Balanced
CPU	
Vendor	Intel
Name	Xeon
Model number	E5-2470
Stepping	C2
Socket type	FCLGA1356
Core frequency (GHz)	2.3
Bus frequency	8 GT/s
L1 cache	32 KB + 32 KB (per core)
L2 cache	256 KB (per core)
L3 cache	20 MB
Platform	
Vendor and model number	Dell PowerEdge M420
Motherboard model number	0MN3VC
BIOS name and version	Phoenix 1.2.4
BIOS settings	Default
Memory module(s)	
Total RAM in system (GB)	96
Vendor and model number	Samsung M393B2G70BH0-YH9
Type	PC3L-10600R
Speed (MHz)	1,333
Speed running in the system (MHz)	1,333
Timing/Latency (tCL-tRCD-tRP-tRASmin)	9-9-9-36

System	Dell PowerEdge M420
Size (GB)	16
Number of RAM module(s)	6
Chip organization	Double-sided
Rank	Dual
Operating system	
Name	Oracle Enterprise Linux 6.2, x86_64
File system	ext4; Oracle ASM volume manager
Kernel	2.6.32-220.el6.x86_64
Language	English
Graphics	
Vendor and model number	Matrox® G200eR
Graphics memory (MB)	16
Driver	2.4.1.0 9/8/2011
RAID controller	
Vendor and model number	Dell PERC H310 Embedded
Firmware version	20.10.1-0084
Hard drive	
Vendor and model number	Dell M16CSD1-50UCV-D
Number of disks in system	2
Size (GB)	50
Buffer size (MB)	N/A
RPM	N/A
Type	SSD
Ethernet adapters	
Vendor and model number	4x Broadcom® BCM57810 NetXtreme® II 10 GigE
Type	Mezzanine
Driver	7.2.8.0 3/13/2012
USB ports	
Number	2 external
Type	2.0

Figure 5: The server configuration we used for this reference architecture.

ABOUT PRINCIPLED TECHNOLOGIES

Principled Technologies, Inc.
1007 Slater Road, Suite 300
Durham, NC, 27703
www.principledtechnologies.com

We provide industry-leading technology assessment and fact-based marketing services. We bring to every assignment extensive experience with and expertise in all aspects of technology testing and analysis, from researching new technologies, to developing new methodologies, to testing with existing and new tools.

When the assessment is complete, we know how to present the results to a broad range of target audiences. We provide our clients with the materials they need, from market-focused data to use in their own collateral to custom sales aids, such as test reports, performance assessments, and white papers. Every document reflects the results of our trusted independent analysis.

We provide customized services that focus on our clients' individual requirements. Whether the technology involves hardware, software, Web sites, or services, we offer the experience, expertise, and tools to help our clients assess how it will fare against its competition, its performance, its market readiness, and its quality and reliability.

Our founders, Mark L. Van Name and Bill Catchings, have worked together in technology assessment for over 20 years. As journalists, they published over a thousand articles on a wide array of technology subjects. They created and led the Ziff-Davis Benchmark Operation, which developed such industry-standard benchmarks as Ziff Davis Media's Winstone and WebBench. They founded and led eTesting Labs, and after the acquisition of that company by Lionbridge Technologies were the head and CTO of VeriTest.