Technical Report

Oracle RAC Built on FlexPod with VMware Solutions Guide

Naveen Harsani, NetApp
January 2013 | TR-4119

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

1 Executive Summary

Oracle® Real Application Cluster (RAC) built on FlexPod® with VMware® is a pretested solution that is suitable for organizations that are planning to consolidate their existing database environment and move to a cloud-based infrastructure in which they can scale up and scale out based on their requirements. This solution is jointly tested and validated by Cisco and NetApp and provides customers with a prescriptive approach to architecting and designing an optimized database solution that minimizes deployment risks and accelerates business results.

1.1 Purpose

This document focuses on the consolidation and scalability capabilities of the Oracle RAC built on FlexPod with VMware solution. This document is not intended as a reference for an end-to-end FlexPod deployment. Refer to the base FlexPod Solution Guide for further reference. The key objectives of this document are to demonstrate the consolidation study and to demonstrate the scalability of the solution in database environments.

1.2 Target Audience and Assumptions

This document is for sales engineers, field consultants, professional services personnel, IT managers, partners, and customers who want to deploy database solutions using the FlexPod solution for VMware architecture.

This document assumes that the reader has:

- NetApp administration skills and administrative access to the storage systems using the command-line interface
- Basic understanding of FlexPod technology and its different configurations
- Basic understanding of Oracle products and technologies

2 Oracle Database and FlexPod Solution Overview

The current industry trends in data centers indicate an increased adoption of virtualization and increased momentum toward a shared IT infrastructure.

To meet these requirements, Cisco and NetApp jointly developed a highly advanced and improved infrastructure solution called FlexPod. FlexPod is an ideal solution for shared, virtualized, and cloud infrastructure requirements.

A typical Oracle Database deployment requires huge system considerations, design, and planning activities. The use of a FlexPod solution reduces these requirements by leveraging pretested, prevalidated, and predictable workloads, thereby enabling a balanced I/O configuration design that is suitable for your environment.

2.1 Problem Statement and Proposed Solution

One of the basic issues faced by CIOs today is managing server and storage sprawl in their data center. This challenge is driving increased growth in financials and labor demands in every industry.

To address this challenge, Cisco and NetApp developed a highly innovative and differentiated solution built on FlexPod technology. This solution uses the latest technology from Cisco, NetApp, and VMware to provide a virtualized infrastructure that helps increase the agility of database environments. FlexPod is a pretested data center solution built to meet the shared infrastructure requirements for virtualized and cloud-based deployments.

FlexPod technology has a defined set of hardware and software components that serve as a flexible, integrated infrastructure stack for all virtualization solutions. The FlexPod for VMware solution includes NetApp® storage, Cisco® networking switches, Cisco Unified Computing System™ (UCS™) x86 blade servers, and VMware vSphere® virtualization software in a single integrated offering. This solution allows the compute, network, and storage components to fit in one data center rack, and it permits expansion to a separate rack as and when required.

2.2 Oracle RAC Built on FlexPod with VMware

Oracle RAC built on FlexPod with VMware is an optimized and balanced solution that can be used to move existing data center practices to a highly flexible and integrated infrastructure. Figure 2 illustrates an Oracle RAC solution that allows databases to be consolidated using the VMware hypervisor. Figure 3 illustrates how a single Cisco chassis can be scaled to meet high-performance and throughput requirements using the scale-up approach of Cisco Unified Computing Systems.

Over the years, Cisco and NetApp have jointly developed a number of customer-centric solutions and use-case documents to help customers transform their data center to a shared infrastructure.

A base FlexPod solution includes the following predesigned and prevalidated components.

- Cisco UCS and Cisco UCS Manager
- Cisco Nexus® switches
- NetApp FAS controllers

Based on performance and reliability requirements, a FlexPod solution provides the flexibility to select one or more of these components to meet the customer's unique solution requirements. Alternatively, customers can choose to start with a base solution as seen in Figure 1 and then increase capacity and performance as and when required.

Figure 1) Base FlexPod solution.

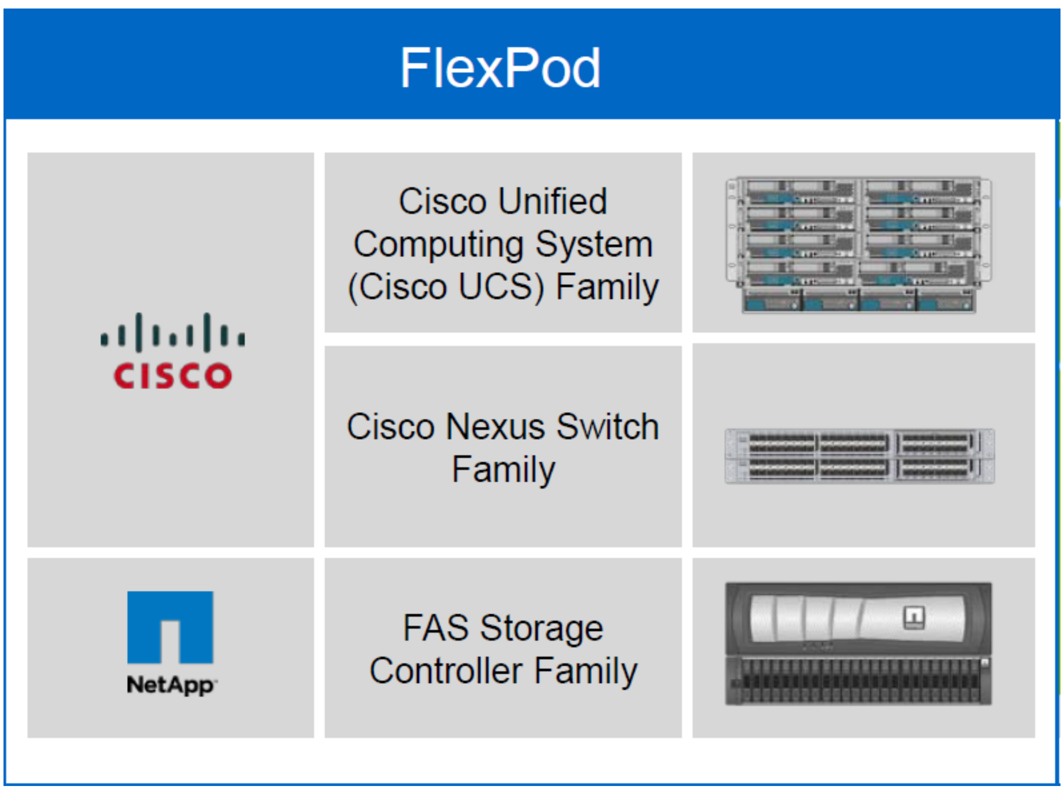

For details on FlexPod, refer to the [FlexPod Data Center Solution] document.

2.3 Typical Use Cases

Oracle RAC built on FlexPod with VMware can be used in several ways; however, this document focuses only on two important use cases:

- Database consolidation and virtualization
- Scalability

Database Consolidation

Customers today want to consolidate their infrastructure to reduce their data center costs and complexity while improving their business agility. In this example, a VMware hypervisor is used to build a virtualized infrastructure to avoid dependency on the physical hardware. A single physical machine can now host numerous VMs, which in turn can host numerous databases and reduce the amount of space that would be required to host all these services.

Database consolidation can also increase effective utilization of the equipment and in turn help reduce the cost of acquiring these licenses to host databases. By using existing licenses more effectively, you can reduce the overall IT expenditure on your infrastructure. For example, one CPU core licensed to use an Oracle Database can host any number of virtual machines, and in turn all of them can virtually host Oracle Databases.

NetApp vFiler® units can help create a secure multi-tenant environment that consolidates numerous business units and functions into one centralized location for easy management and security.

The other important benefit of hosting a database in a virtualized infrastructure is additional agility. VMs can be deployed and reprovisioned to create additional VMs without affecting the existing VMs. Customers can use NetApp FlexClone® technology to deploy numerous VMs without actually creating physical copies of the data blocks. This helps reduce the storage footprint and in turn reduces the cost required to acquire additional disk space.

Figure 2) Database consolidation.

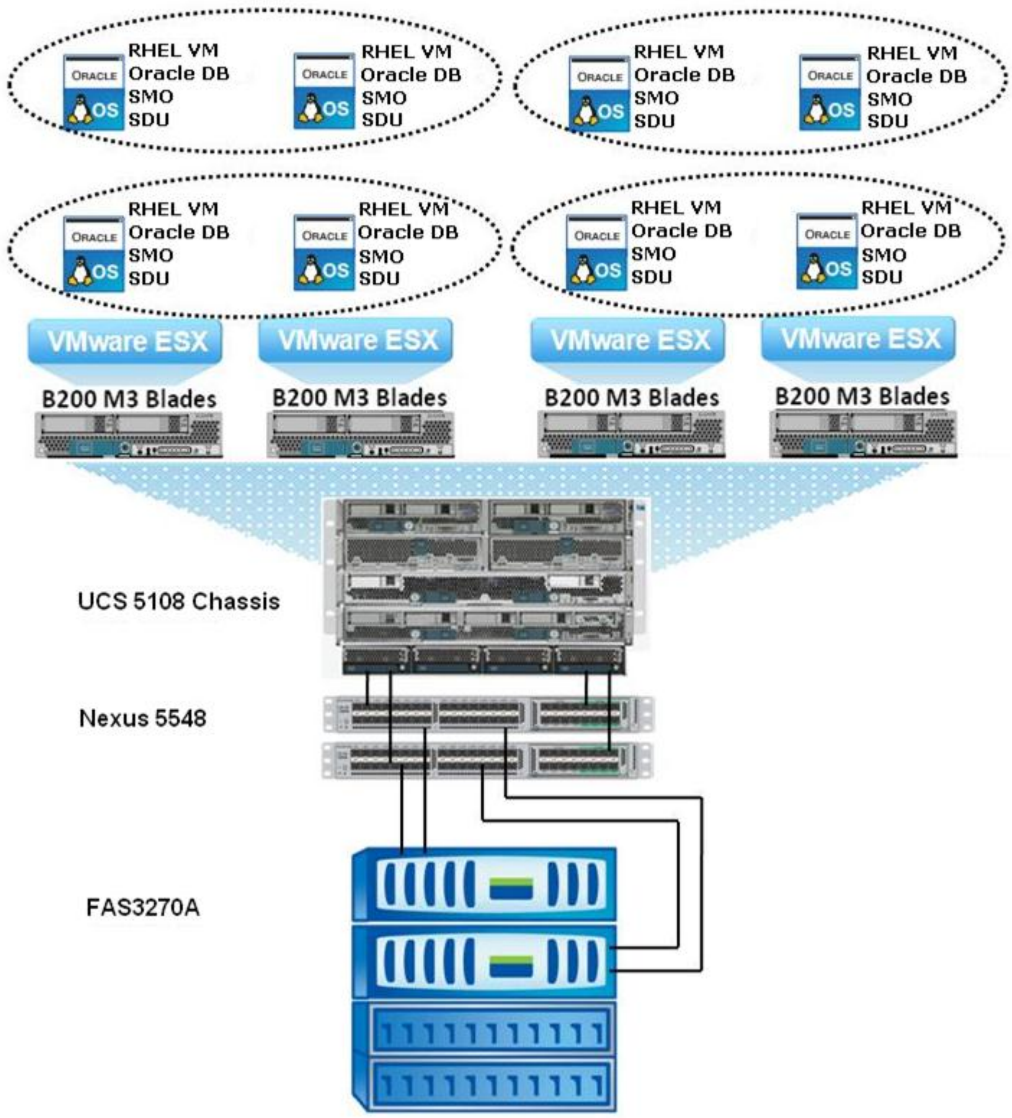

Note: Abbreviations used in Figure 2 are:

RHEL VM—Red Hat Enterprise Linux® Virtual Machine

Oracle DB—Oracle Database

SMO—SnapManager® for Oracle

SDU—SnapDrive® for UNIX®

Improved Scalability

This section defines scalability and how it is enabled in terms of Oracle in a cloud architecture. A scalable architecture is usually misunderstood as the number of cores or the scale of disk drives that can be populated in a system. However, scalability is not just limited to this alone.

A scalable system is one that is designed to scale as required while maintaining necessary service levels and required resource utilization. Usually in a complex environment it is very difficult to maintain and monitor the resource utilization, along with keeping track of the SLAs.

FlexPod helps to create a balanced configuration for all compute, network, and storage layers to justify an optimized and scalable data center solution. The VMware hypervisor helps to provision only the required amount of memory and processing power to each VM, which can then be scaled up to a desired level. NetApp Flash Cache™ intelligent caching helps to improve read throughput by caching the frequently accessed data. This feature is built in to the NetApp controllers and helps to reduce average latency, which in turn reduces the overall transaction response time. Several users can now simultaneously work on the system.

Figure 3) Improved scalability.

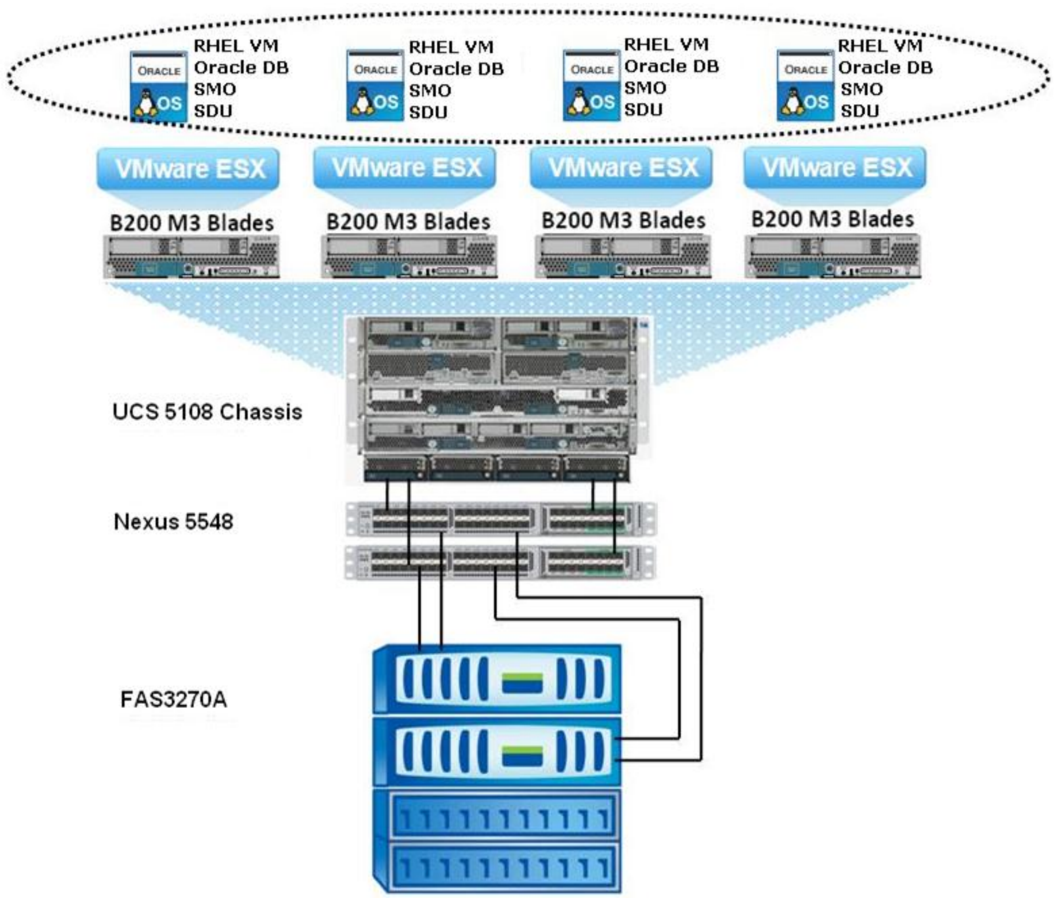

3 Technology Overview

FlexPod is a flexible predesigned infrastructure configuration that is built on the Cisco UCS, Cisco Nexus switches, and NetApp FAS storage. VMware is used on top of FlexPod technology to build a virtualized infrastructure. This infrastructure also provides the basis for a private cloud environment setup. FlexPod can scale up for greater performance and capacity, and it can scale out for environments that need consistent, numerous deployments. FlexPod has a baseline configuration, but it also has the flexibility to be sized and optimized to accommodate many different business solutions and use cases.

Cisco and NetApp have provided well-defined documentation for best practices and the deployment collateral necessary to build the FlexPod with VMware shared infrastructure stack. As part of the FlexPod offering, Cisco and NetApp have created a white paper with a technical specification sheet, "FlexPod

Technical Specifications," that is highly modular. Although each FlexPod unit might vary in its configuration to meet the customer's initial use case, once a FlexPod unit is built it can easily be scaled as requirements and demands change. This includes both scaling up (adding additional resources within a FlexPod unit) and out (adding additional FlexPod units).

4 Design Considerations

FlexPod is defined with a set of hardware and software that serves as a foundation for data center deployment that includes NetApp storage, Cisco networking, and Cisco UCS in a single package. The solution can be scaled while still maintaining its integrity by augmenting a single FlexPod instance to provide the appropriate network, compute, or storage capacity needed within a single pod. It can also be scaled by using the port density of the Cisco Nexus 5500 networking platform to accommodate numerous instances of FlexPod. In either case, the flexibility of the pod construct allows numerous solutions to be built on top of one or more FlexPod configurations, providing enterprise flexibility, supportability, and manageability for both virtualized and nonvirtualized environments. Figure 2 illustrates a few FlexPod scaling options. Figure 1 illustrates a base FlexPod configuration with hardware families and solution benefits.

The NetApp Data ONTAP® 8.0.1 architecture provides support for 64-bit aggregates that support larger aggregate sizes compared to the current 32-bit limit of 16TB. FlexVol® volumes created with 64-bit aggregates also support much larger volumes ranging from 30TB to 100TB in size, based on the storage system. Refer to the following documents for more information.

- The system configuration guides for maximum supported 8.0.1 volume and aggregates sizes for specific hardware platforms on the NetApp Support site.
- TR-3786: A Thorough Introduction to 64-Bit Aggregates

4.1 Flash Cache

A Flash Cache module can optimize the performance of your random-read-intensive workloads such as online transaction processing, file services, and messaging. Flash Cache works with both 32-bit and 64-bit aggregates and caches data that comes from volumes located in both types of aggregates. Flash Cache caches data based on the data-access pattern regardless of the aggregate type.

The data cached in Flash Cache while the system is in operation depends on the workload, and it can be a combination of data from volumes contained in different aggregates. Sixty-four-bit aggregates have a bigger address space and also use more memory for their metadata compared to 32-bit aggregates. Flash Cache complements this by providing caching for both data as well as metadata to improve the overall response time.

4.2 NetApp Storage Configuration Overview

This section discusses the NetApp storage layout design considerations required when deploying an Oracle Database 11*g*™ R2 GRID infrastructure with RAC option on a VMware ESXi™ 5.0 hypervisor on Cisco UCS in an NFS network environment.

Figure 4 illustrates a high-level storage design on a NetApp FAS3270 HA storage system.

Figure 4) High-level storage design.

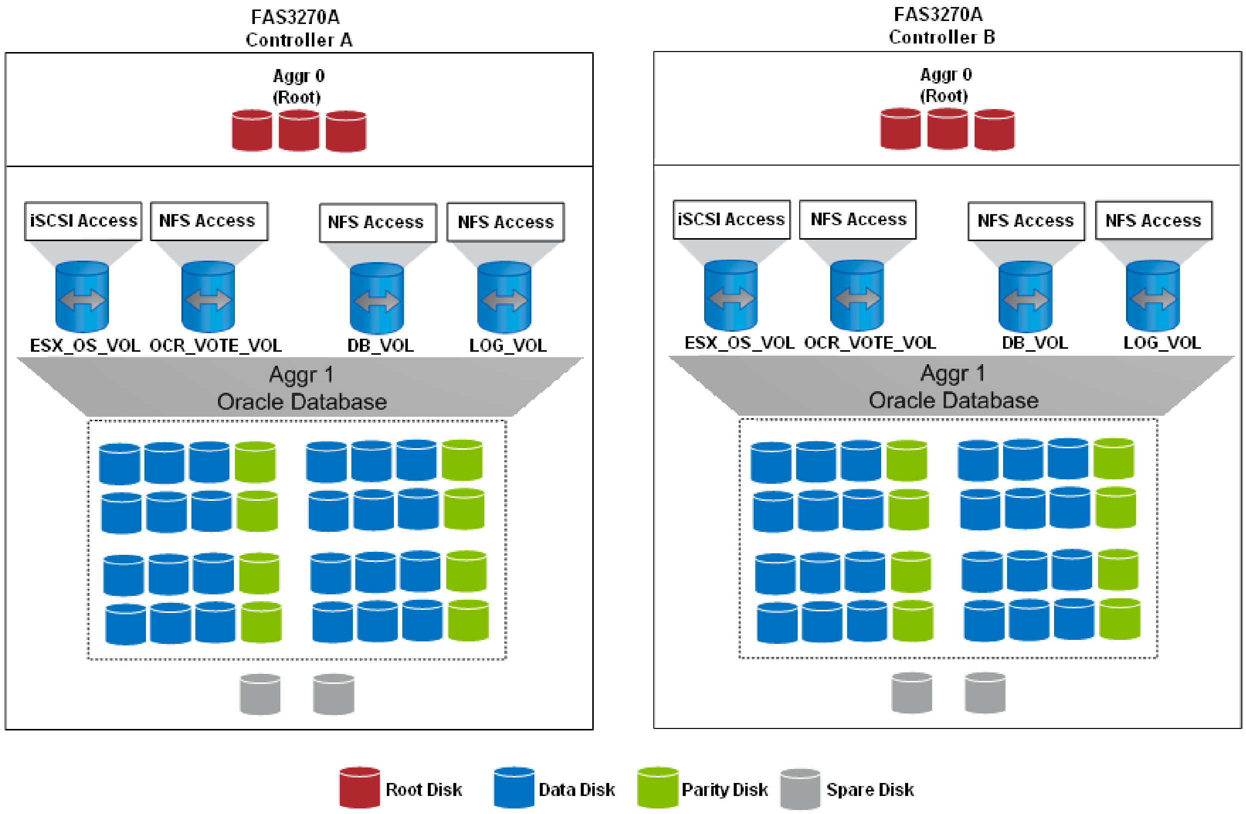

Table 1 lists the NetApp storage layout.

Table 1) NetApp storage layout.

Aggregation and NetApp Controller	NetApp FlexVol	Flexible LUN	Comments
Aggr1 on Controller A	Boot_OS_VOL	ESXi_OS_LUN	iSCSI boot LUN for VMware ESXi host for node 1 and node 2 of failover cluster, with Cisco UCS B200M3 blade server
Aggr1 on Controller A	OCR_VOTE_VOL		Used to store OCR and voting disk by using NFS
Aggr1 on Controller A	DB_VOL		Used to store data files, spfile, and a copy of control files
Aggr1 on Controller A	LOG_VOL		Used to store redo log files and a copy of control files
Aggr1 on Controller B	Boot_OS_VOL	ESXi_OS_LUN	iSCSI boot LUN for VMware ESXi host for node 3 and node 4 of failover cluster, with a Cisco UCS B200M3 blade server
Aggr1 on Controller B	OCR_VOTE_VOL		Used to store OCR and voting disk by using NFS

Aggregation and NetApp Controller	NetApp FlexVol	Flexible LUN	Comments
Aggr1 on Controller B	DB_VOL		Used to store data files, spfile, and a copy of control files
Aggr1 on Controller B	LOG_VOL		Used to store redo log files and a copy of control files

Follow these steps to create additional 64-bit aggregates, determine the aggregate name and the storage controller on which to create it, and how many disks it will contain.

1. Execute the following command to create a new aggregate:

```
aggr create <<var_aggr01>> -B 64 <<var_num_disks>>
```

Note: Leave at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

2. To configure the Network File System (NFS) service on a storage system, complete the following steps:
   a. Add a license for NFS.

```
license add <<var_nfs_license>>
```

   b. Set the following recommended options that enable NFS version 3 (NFSv3).

```
options nfs.tcp.enable on
options nfs.udp.enable off
options nfs.v3.enable on
```

   c. Enable NFS.

```
nfs on
```

3. Run these steps on both storage controllers in a high-availability (HA) pair.

4.3 NetApp FAS3270 HA

Use the following commands to configure NetApp storage systems to implement the storage layout design.

1. The following command creates Aggr1 with a RAID group size of 16, 55 disks, and RAID_DP redundancy for hosting NetApp FlexVol volumes and LUNs as shown in Table 7.

```
FAS3270HA-Controller A> aggr create aggr1 -t raid_dp -r 16 55-B 64
```

2. The following commands create NetApp FlexVol volumes on Aggr1 for hosting iSCSI LUNs and database volumes as described in Table 1. These volumes are exposed to VMware ESXi host and guest virtual machines.

```
FAS3270HA-Controller A> vol create Boot_OS_VOL aggr1 200g
FAS3270HA-Controller A> vol DB_VOL aggr1 1024g
FAS3270HA-Controller A> vol create LOG_VOL aggr1 500g
FAS3270HA-Controller A> vol create OCR_VOTE_VOL aggr1 10g
```

3. Repeat the aggregate and FlexVol creation steps on Controller B after creation of FlexVol volumes on Controller A.
4. The following command creates LUNs on FlexVol volumes for iSCSI boot of ESXi host.

```
FAS3270HA-Controller A> lun create -s 150g -t vmware /vol/Boot_OS_VOL/ESXi_OS_LUN
```

5. Repeat the LUN creation steps on Controller B after creation of the LUN on Controller A.
6. The following command creates an initiator group (igroup) for mapping the VMware ESXi host boot LUN.

```
FAS3270HA-Controller A> igroup create –I –t vmware iSCSI-ESXi –Boot iqn.2012-01.com.vmware:ESXi
```

7. Repeat step 6 on Controller B after creation of the igroup on Controller A.
8. The following command maps LUNs to a specific igroup to access the VMware ESXi host boot.

```
FAS3270HA-Controller A>
lun map /vol/Boot_OS_VOL/ESXi_OS_LUN  iSCSI-ESXi-Boot
```

9. Repeat the preceding commands on Controller B after they are completed on Controller A.

NFS exports all the flexible volumes (data volumes, redo log volumes, and OCR/voting disk volumes) from both Controller A and Controller B, providing read/write access to the root user of all hosts created in the previous steps.

5 Oracle RAC Deployment

Oracle RAC is available for Oracle Database software beginning with version 9*i*™. The key feature of Oracle RAC is that it can provide high-availability and clustering capability to a multinode server farm.

Oracle RAC allows numerous computers to run Oracle RDBMS software while simultaneously accessing a single database. The concurrent access to the same database is provided by a cluster stack called Oracle Clusterware. The single database is managed by two more clustered nodes by using Cache Fusion technology. Cache Fusion allows fusing of the in-memory data that is cached physically and separately on each computer into a single global cache.

RAC deployment can be done in many different ways. Therefore, it is important to choose one of the installation types and then proceed based on the requirements. This solution guide focuses only on Direct NFS deployment. However, if you want to deploy RAC in other configurations, refer to the NetApp Interoperability Matrix to determine the supported configurations.

Table 2 describes the production hardware resources required for deployment.

Table 2) Production hardware resources.

Local Directory on Guest OS	NetApp NFS Volumes	Owner	Purpose
/u01/app/11.2.0/grid	NA	Grid	Oracle GRID binary installation
/u01/app/oracle	NA	Oracle	Oracle Database binary installation
/data_A	/vol/RAC1_DB	Oracle	Data files and control files
/data_B	/vol/RAC1_DB_B	Oracle	Data files and control files
/log_A	/vol/RAC1_LOG	Oracle	Redo log files and control files
/log_B	/vol/RAC1_LOG_B	Oracle	Redo log files and control files
/ocrvote	/vol/OCR_VOTE1	Grid	OCR and voting disks

1. Edit /etc/fstab file in each RAC node and add the entry for all volumes and its corresponding local directories created in the preceding steps with the proper mount options, as illustrated in Figure 5.

Figure 5) Oracle mount options.

```
192.191.1.2:/vol/OCR_VOTE1      /ocrvote        nfs rw,bg,hard,rsize=65536,wsize=65536,vers=3,actimeo=0,nointr,suid,timeo=60
0,tcp       0 0
192.191.1.2:/vol/RAC1_DB        /data_A         nfs rw,bg,hard,rsize=65536,wsize=65536,vers=3,nointr,suid,timeo=600,tcp 0 0

192.191.1.4:/vol/RAC1_DB_B      /data_B         nfs rw,bg,hard,rsize=65536,wsize=65536,vers=3,nointr,suid,timeo=600,tcp 0 0

192.191.1.2:/vol/RAC1_LOG       /log_A          nfs rw,bg,hard,rsize=65536,wsize=65536,vers=3,nointr,suid,timeo=600,tcp 0 0

192.191.1.4:/vol/RAC1_LOG_B     /log_B          nfs rw,bg,hard,rsize=65536,wsize=65536,vers=3,nointr,suid,timeo=600,tcp 0 0
```

To find the appropriate mount options for different file systems of Oracle 11*g* R2, refer to [https://kb.netapp.com/support/index?page=content&id=3010189.](https://kb.netapp.com/support/index?page=content&id=3010189)

Note: rsize and wsize of 65536 is supported by NFS v3 and is used in this configuration to improve performance.

2. When the editing of `/etc/fstab` in each RAC node using root user is completed, mount all the local directories created to store database, OCR, and voting disks by using the following command:

```
node1# mount /ocrvote
node1# mount /data_A
node1# mount /data_B
node1# mount /log_A
node1# mount /log_B
```

Every release of an Oracle Database requires the completion of certain prerequisites before installation can begin. NetApp advises customers to refer to the Oracle documentation for details on how to complete these prerequisites. This solution guide assumes that the user has a fair understanding of how to install Oracle in an operating environment. Creating users and directories, assigning permissions, and mounting the file system are covered under the Prerequisites section of the Oracle Database Installation Guide.

5.1 Installing the Oracle Grid Infrastructure on Linux

To install the Oracle Grid infrastructure on Linux, complete the following steps.

1. Download the Oracle Grid infrastructure image from [http://edelivery.oracle.com/linux.](http://edelivery.oracle.com/linux)
2. Unzip the Oracle Grid infrastructure binaries to the temporary installation directory on the first server, usually called the master node. The installation will begin in a multinode RAC environment.
3. Use an xterm session to log in to the master node as grid user.
4. Run the Oracle Universal Installer from the installation directory selected in step 2.
   a. Enter the My Oracle Support (MOS) credentials to check for updates.
   b. Select Skip Software Updates.
   c. Click Next.
5. On the Grid Infrastructure page select Install and Configure Oracle Grid Infrastructure for a Cluster and click Next.
6. On the Install Option page, select Advanced Installation and click Next.
7. On the Grid Plug and Play Information page, complete the following tasks.
   a. Enter the cluster name **onetb-cluster**.
   b. Enter the SCAN name **SCAN-ONETB**.
   c. Enter the SCAN port **1521**.

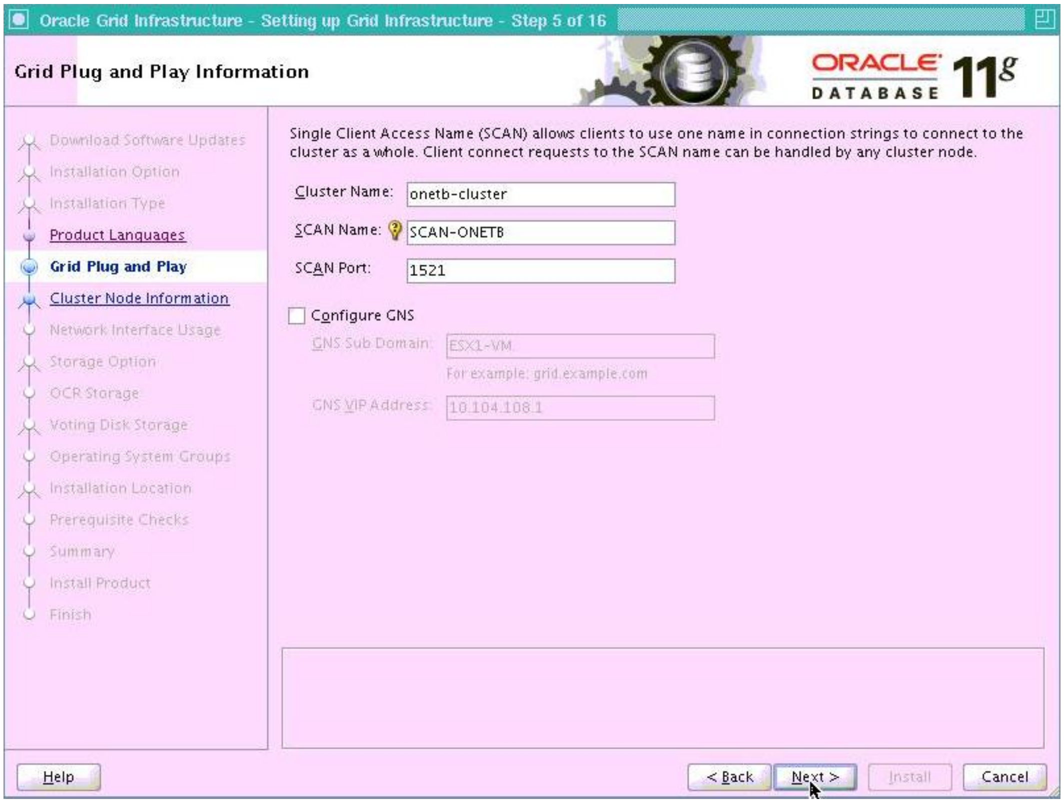

8. Complete the following steps if using the GNS subdomain configuration.
   a. Enter the GNS subdomain var_ORA_GNSSubDomain.
   b. Enter the GNS VIP address var_ORA_GNSVIPAddress.
   c. Click Next.

   For this installation the Configure GNS option has been unchecked.
9. On the Cluster Node Information page, click Add to configure the following node information on the Oracle RAC cluster.
   a. Enter public host name 1 ESX1-VM1.
   b. Enter public host name 2 ESX1-VM2.
   c. Enter virtual host name 1 ESX1-VM1-vip.
   d. Enter virtual host name 2 ESX1-VM2-vip.
   e. Click OK and click SSH Connectivity.
   f. Enter the grid user name for SSH connectivity var_ORA_GridUsername.
   g. Enter the grid password for SSH connectivity var_ORA_GridUserPasswd. SSH connectivity between the nodes is now complete.
   h. Click OK.

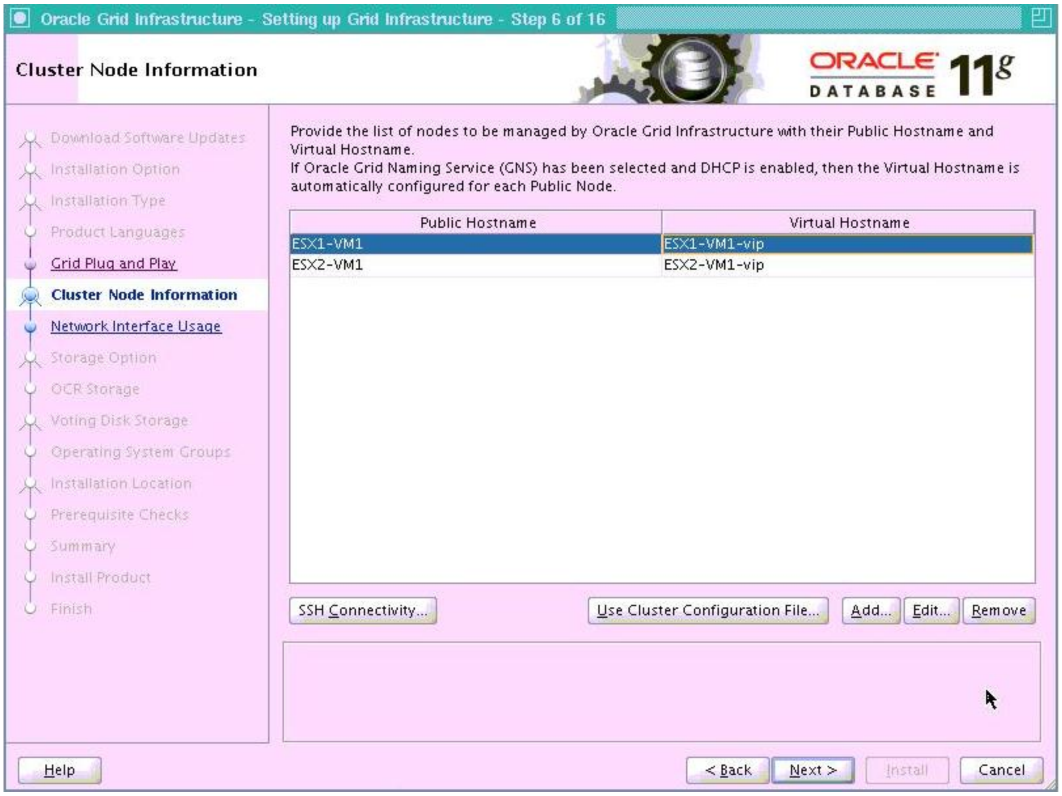

10. On the Specify Network Interface Usage page, verify the network interface names and subnet. Select the appropriate interface type for public and private networks. Click Next.

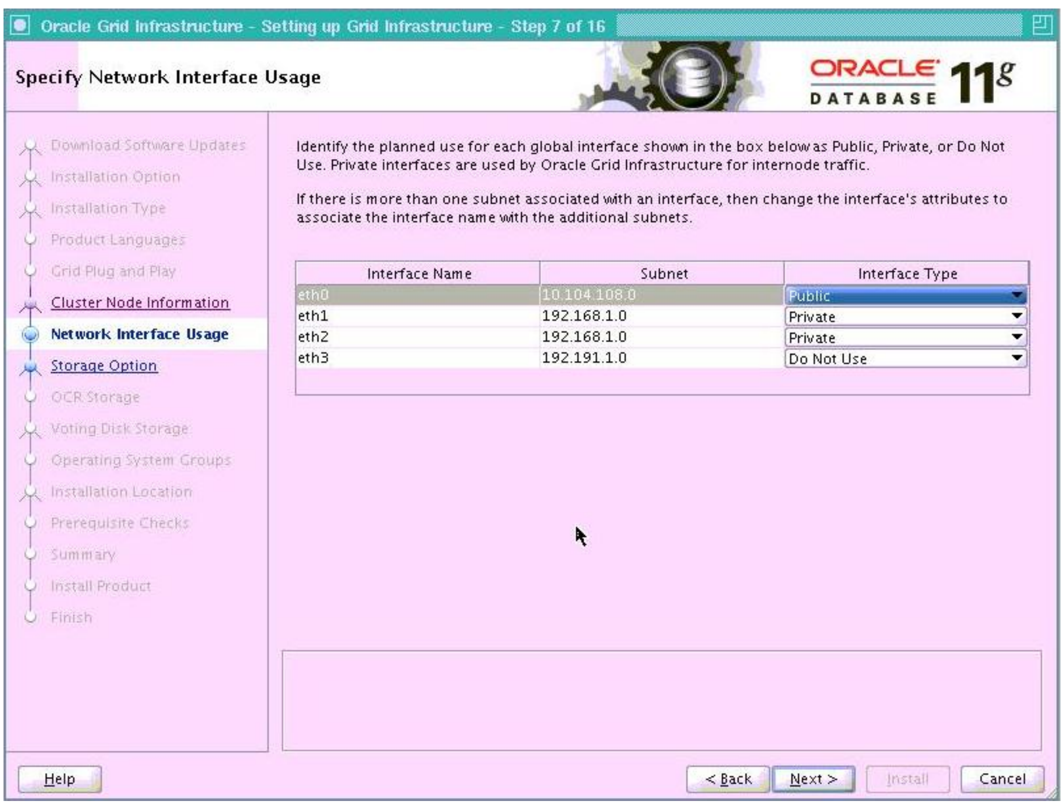

11. On the Storage Option page, select Oracle ASM and click Next.
12. On the Create ASM Disk Group page for creating disk groups (Oracle Cluster Registry [OCR] and Vote), review the disk paths and select the candidate disks to use for clusterware. Select High Redundancy and click Next.

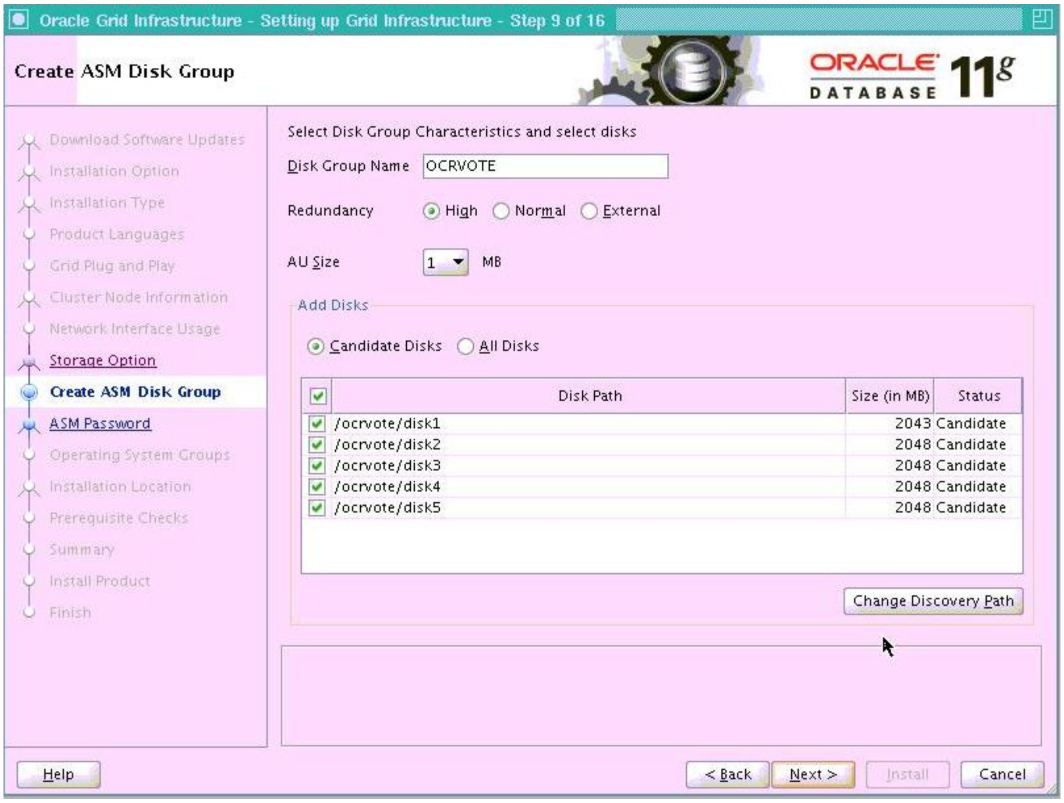

Note: Refer to the Oracle ASM section for configuring and creating ASM disks with the LUNs.

13. On the ASM Password Management page, complete the following tasks.
    a. Select Use Same Accounts for these accounts.
    b. Enter the password `var_ORA_ASM_password` in the Password field.
    c. Reenter the same password `var_ORA_ASM_password` in the Confirm Password field.
    d. Click Next.
14. On the Operating System Groups page, complete the following tasks.
    a. Select `dba` from the Oracle ASM DBA group drop-down list.
    b. Select `dba` from the ASM Operator group drop-down list.
15. This step is optional.
    a. Select `dba` from the Oracle ASM Administrator group drop-down list.
    b. Click Next.
16. On the Specify Installation Location page, verify that the populated Oracle Base and Software Location paths are correct, based on the details picked from the bash profile. Click Next to validate the information that is displayed, after which the OUI continues.

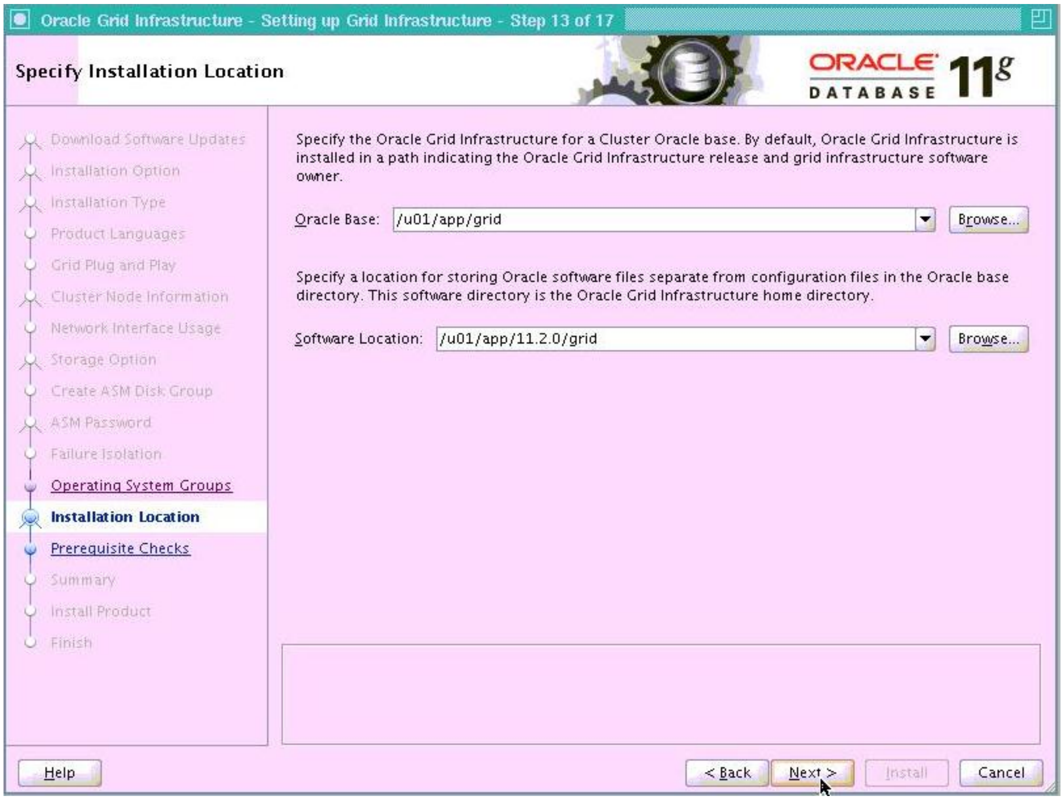

17. On the Prerequisite Checks page, verify that all prerequisite parameters are correct. Click Next. The comprehensive Oracle RAC prerequisites are checked and the failures are displayed.

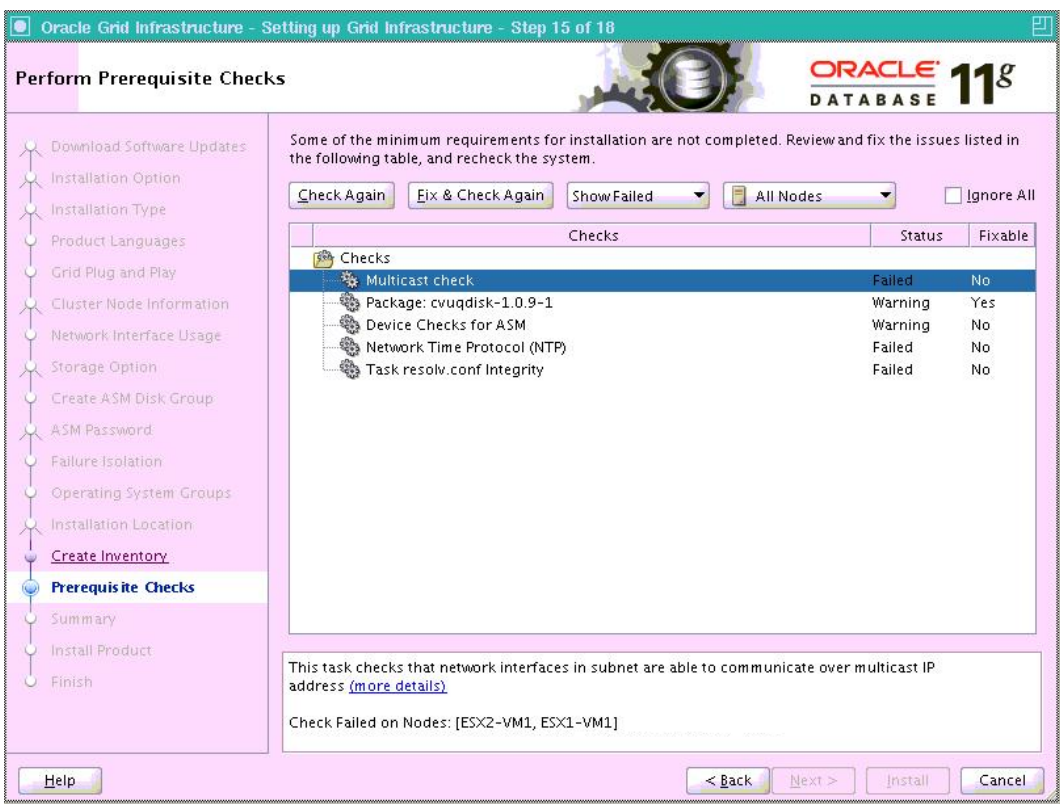

The Execute Fixup Scripts page displays the failure details.

18. Run the scripts shown in the pop-up window to fix the issues on all the nodes, as required. Click OK.

19. On the Execute Configuration Scripts page, run the configuration scripts on all nodes in the cluster. Click OK.

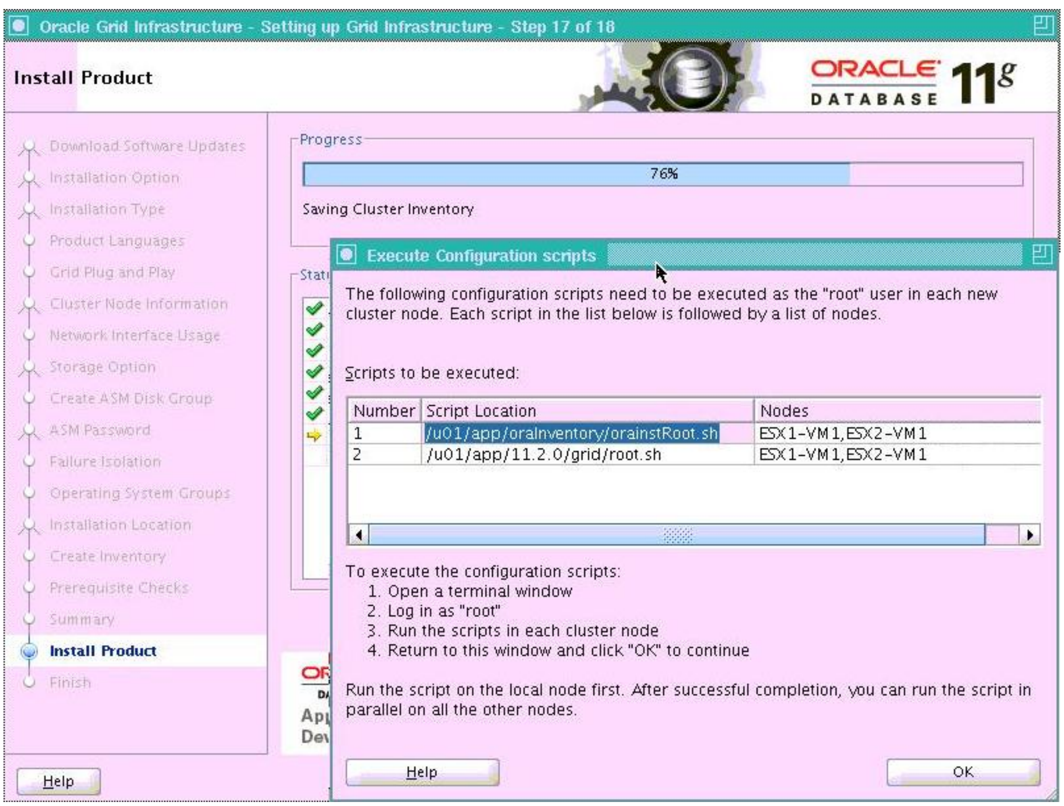

The agent configuration tasks are displayed. After all agent assistants complete successfully, the grid infrastructure installation is complete.

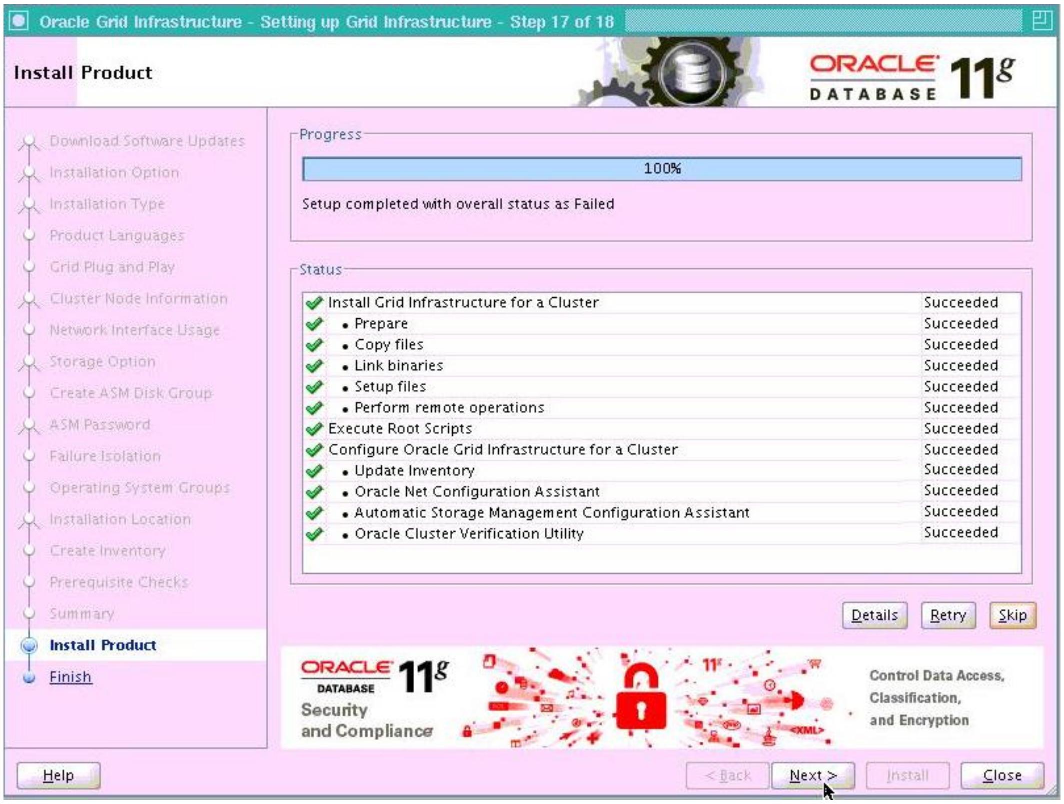

Make sure that sufficient disk space is available for installing the database software and temporary space for temporary file management. For information about the Oracle documentation refer to “Installing Oracle Real Applications Clusters on Linux.”

5.2 Installing Oracle RAC RDBMS

To install Oracle RAC RDBMS, complete the following steps.

1. Download the Oracle Database software from the Oracle Software Delivery Cloud. Use different users for the RAC and Grid Infrastructure installations to achieve job role separation.
2. Check the status of the cluster that is already installed by using the Grid infrastructure.
    a. `<grid_home>/bin crsctl check crs`
    b. CRS-4638: Oracle High Availability Services is online
    c. CRS-4535: Cannot communicate with Cluster Ready Services
    d. CRS-4529: Cluster Synchronization Services is online
    e. CRS-4533: Event Manager is online
3. Browse to the directory where the software is unzipped. This directory is the staging area and contains the Oracle Database distribution.
4. Start the installation by using the runInstaller binary file. The Oracle Universal Installer loads the setup driver and performs the installation prerequisites.

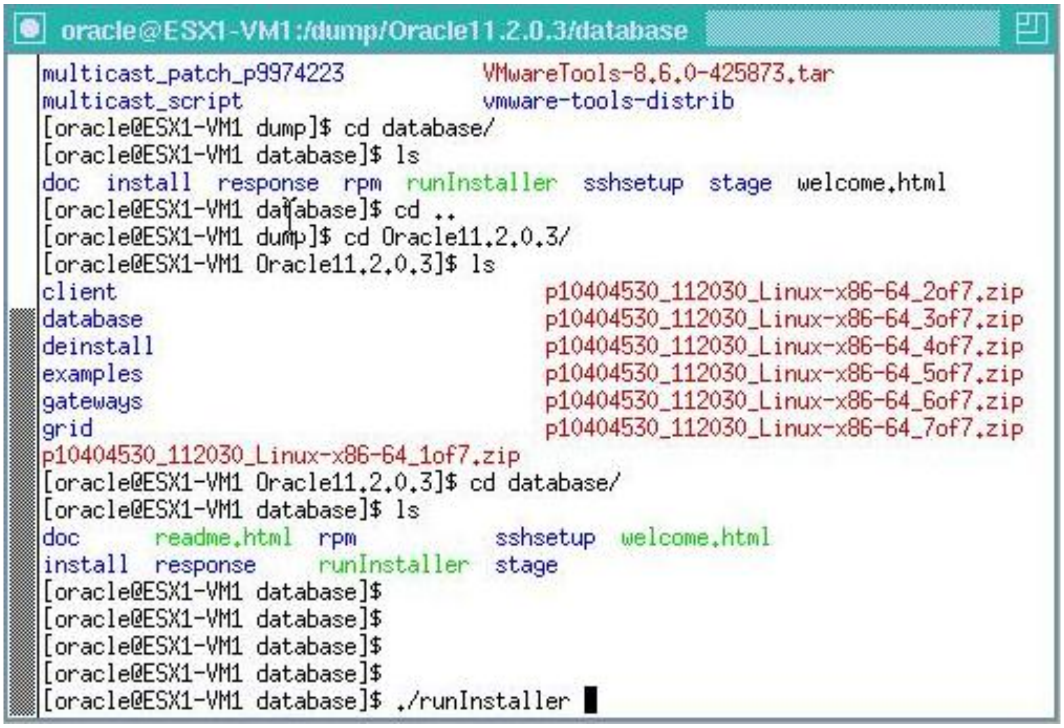

5. Configure security updates. Make sure that proactive support is configured and working and that e-mail notifications are also configured. Click Next.

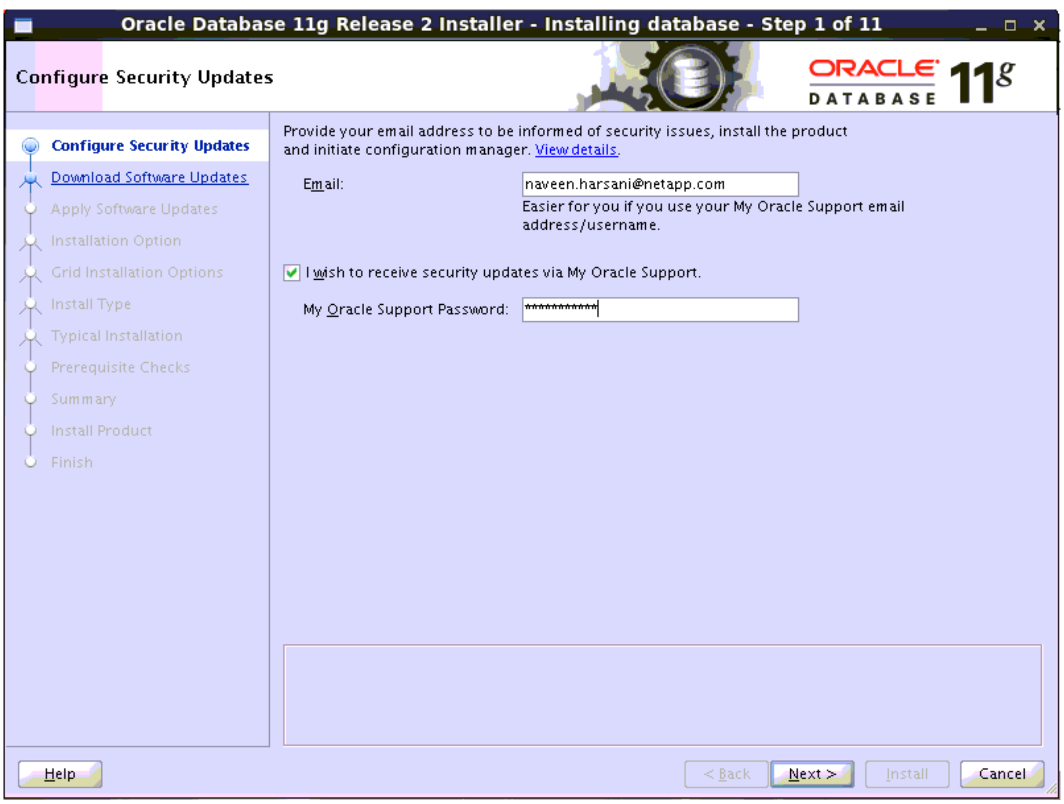

6. Select an installation option. The default preselected option is Create and Configure a Database. NetApp recommends that you select Install Database Software Only for a RAC database installation.

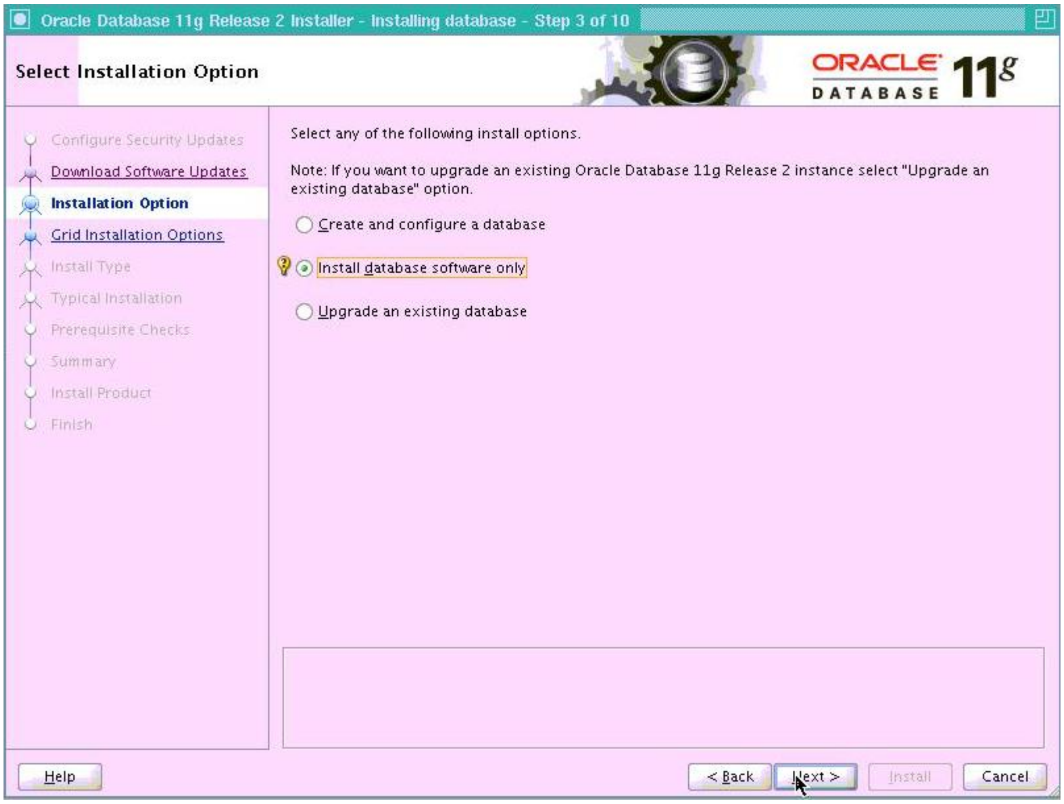

7. Select a grid option. You can perform a single-instance installation or a grid installation across the nodes of the cluster. Select the Real Application Cluster Database Installation option. All selected nodes are automatically installed with the database software.

   Oracle recommends that you select only the nodes in which you want the installation to complete automatically.

8. After you select these nodes, verify whether SSH connectivity is properly configured. If it is not, specify a user name and password by clicking SSH Connectivity and then clicking Setup.

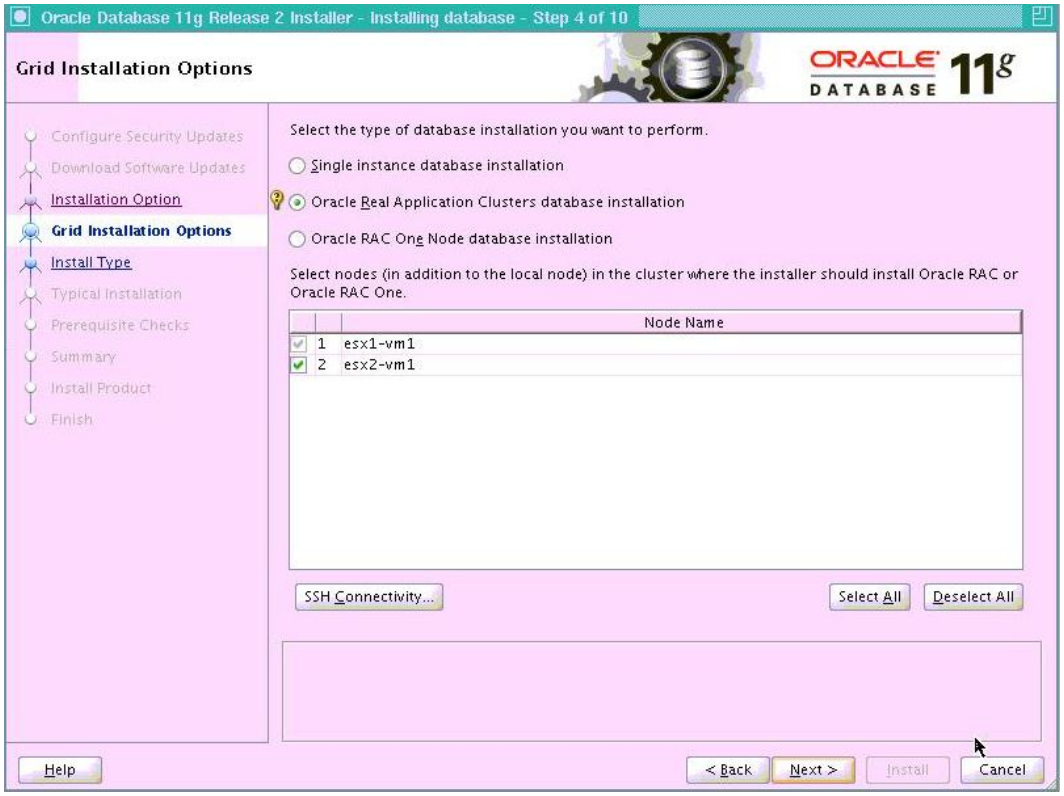

9. Select the languages in which the database will run. The default language for the installation is English.

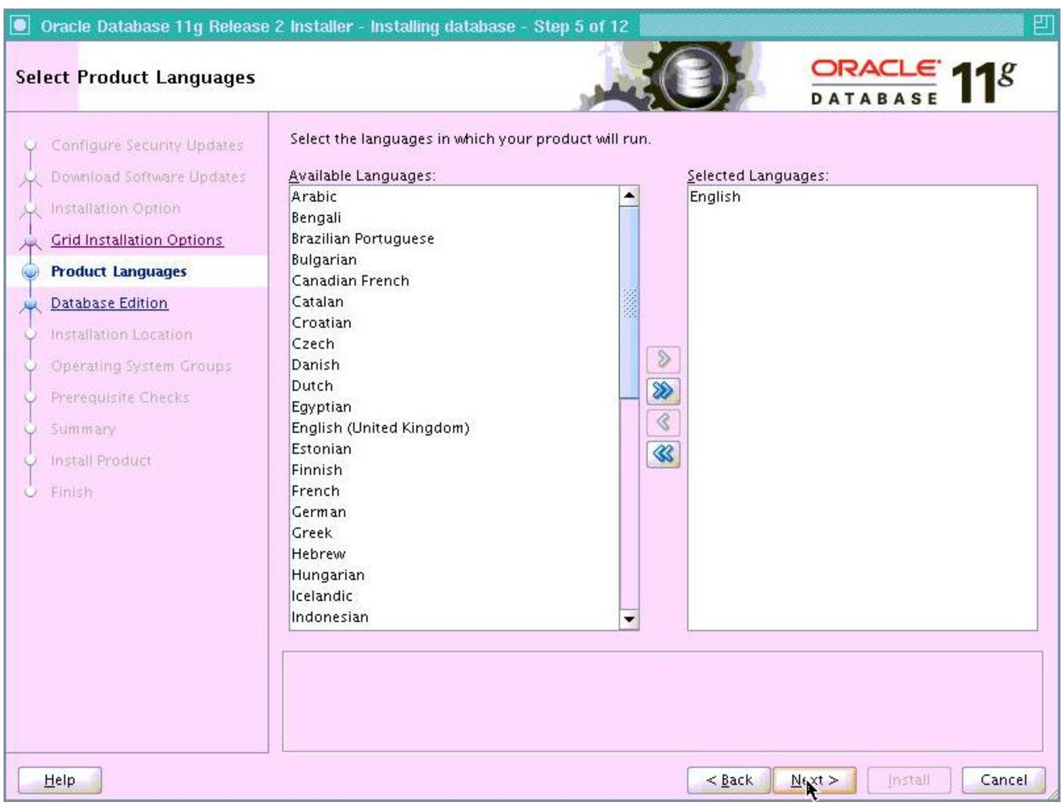

10. Select a database edition to install. In this example, the Enterprise Edition is installed.

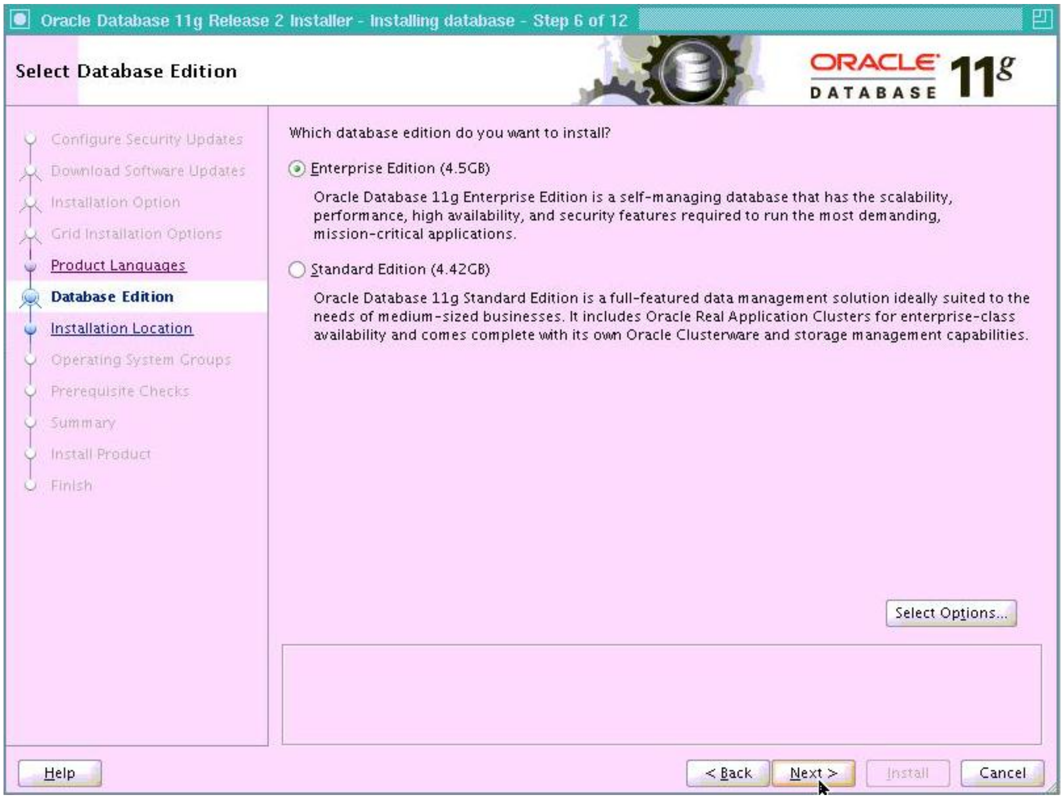

You may also select licenses purchased from Oracle. The default options are preselected. Oracle recommends that you select only the options for which you have a license.

11. Select the location to install Oracle binaries. This location should be a shared location that is available to all cluster nodes.

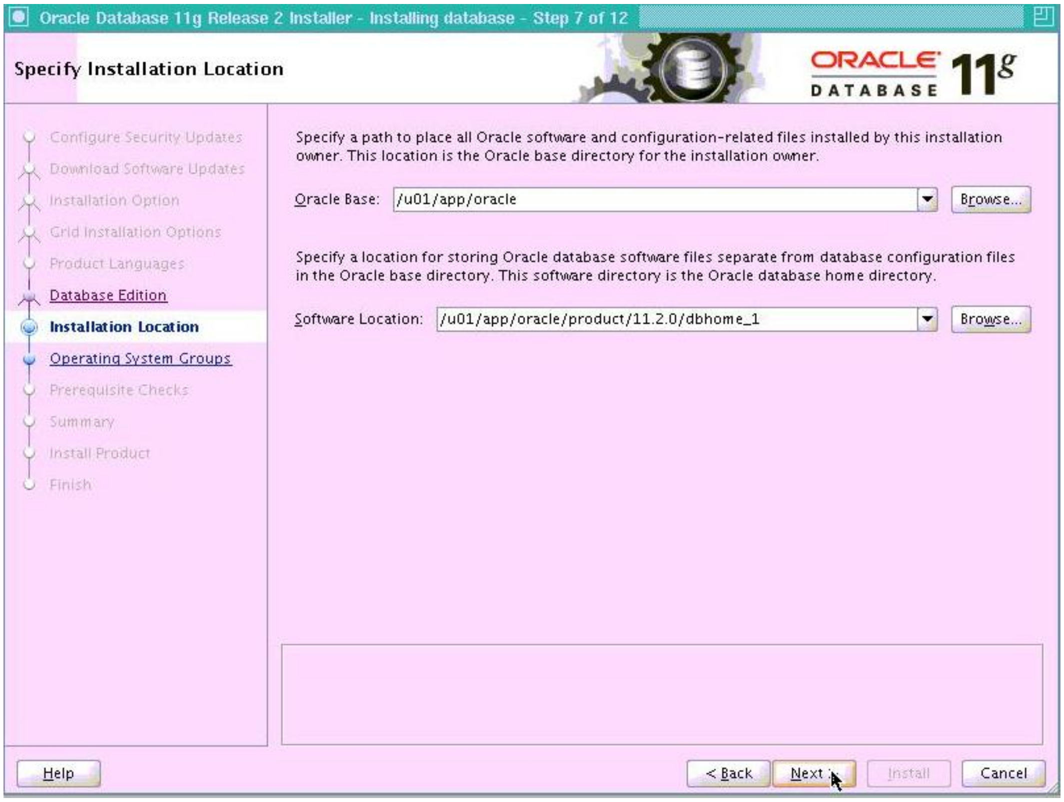

Best Practice
This example uses the same local directory across all nodes to store Oracle binaries. However, a shared Oracle home can be specified for all nodes in a cluster. When a separate local directory installation is selected, rolling updates and patches can be easily applied. When a shared home installation is selected, better storage efficiency can be achieved, but at the cost of availability. This example does not cover an installation based on a shared Oracle home. Rolling upgrades and patching are not allowed when a shared Oracle home is used.

12. Select the OSDBA and OSOPER groups for the installation. If the operating system groups were configured, the installer automatically selects the appropriate groups.

13. Perform the prerequisite system checks to make sure that the installation can proceed.

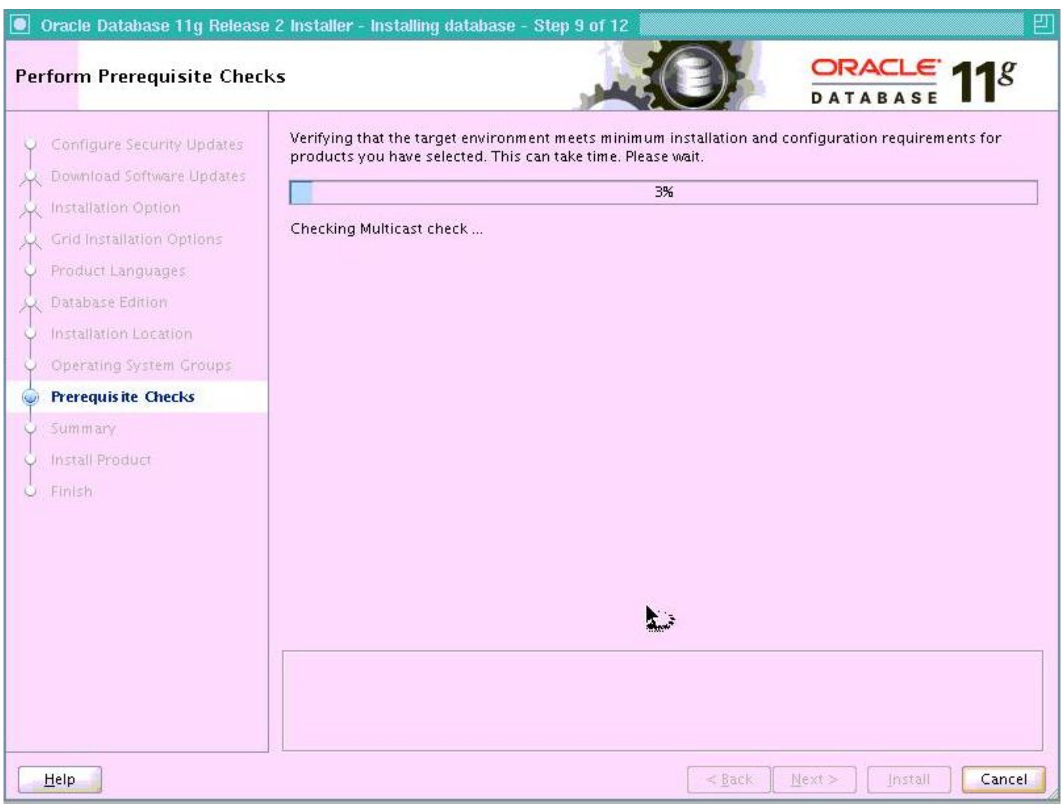

After all checks have successfully completed, the Summary page displays all the options that were selected during the installation procedure.

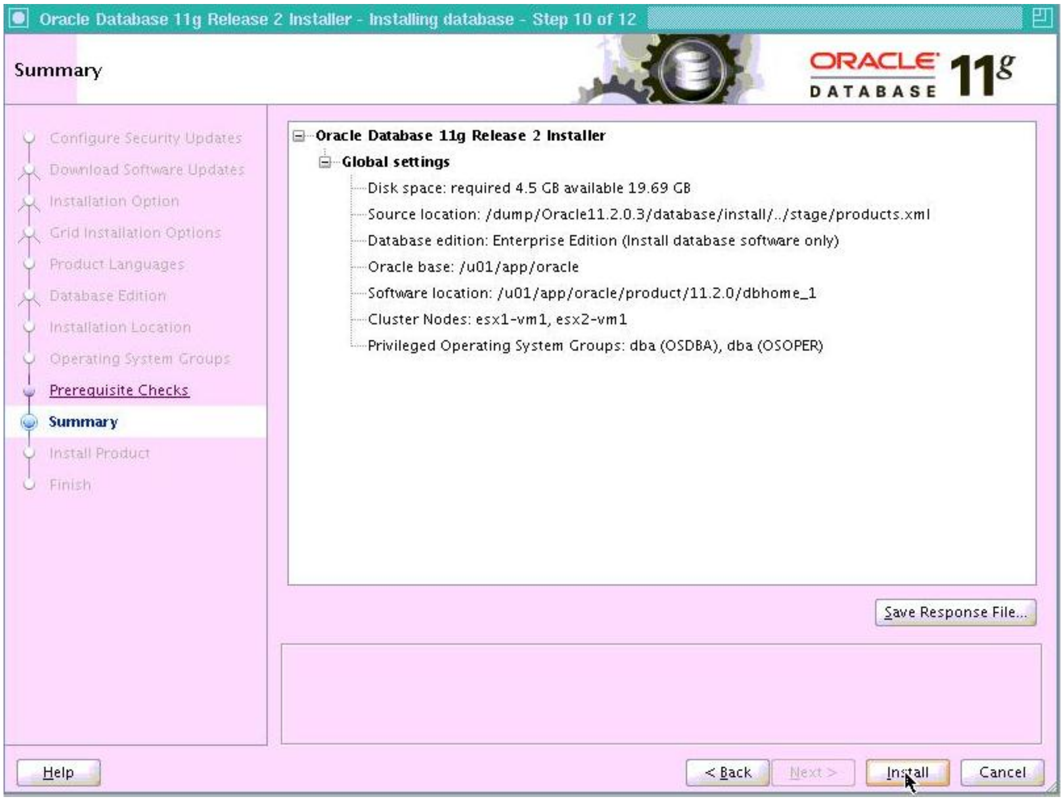

14. Click Install.

    The Oracle installation begins on the node where the installation was initiated and then continues with the other nodes. Follow the steps described in the Status pane.

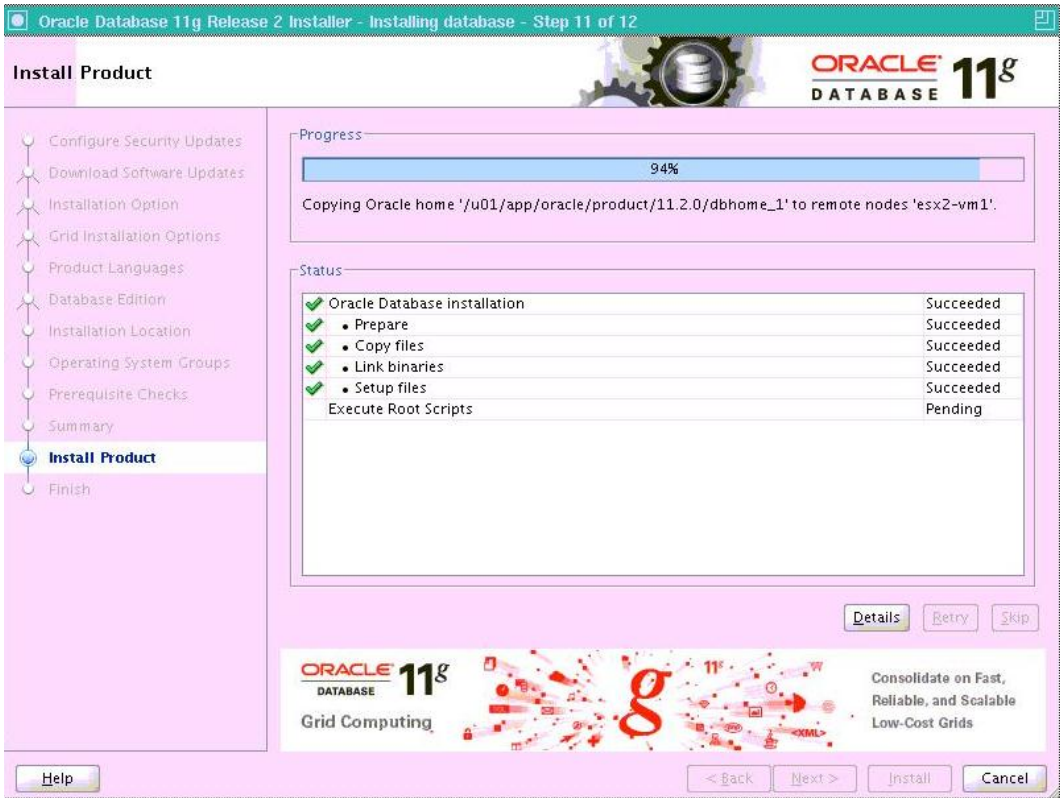

The installation continues on its own. No further input is required until the root script stage described in the following steps.

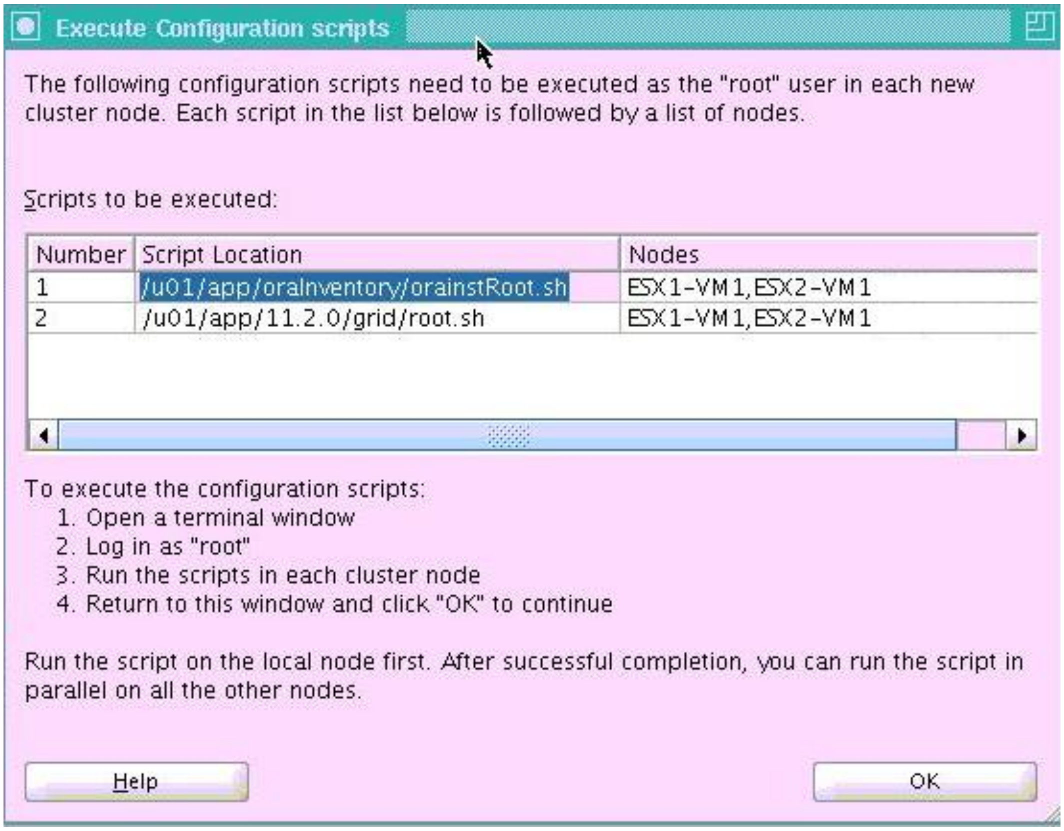

a. To complete the installation, the root user must run a root script on each cluster node. To perform this task, open a new terminal window and run the root script by using the complete path illustrated in the following example. After the `root.sh` has been specified with the complete path, press Enter to run the script.

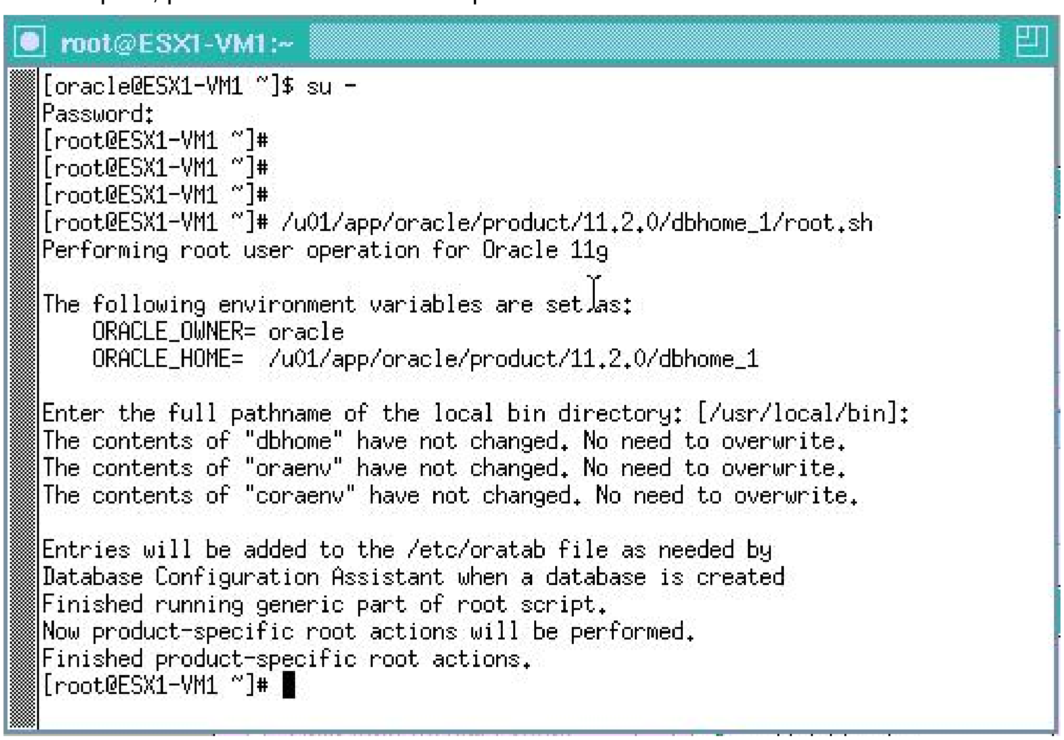

b. Run the same script on the other nodes of the cluster. When finished, click OK in the Execute Configuration Scripts dialog box. Oracle recommends that you allow the installation to complete on its own. Do not interrupt the installation by clicking the Cancel button.

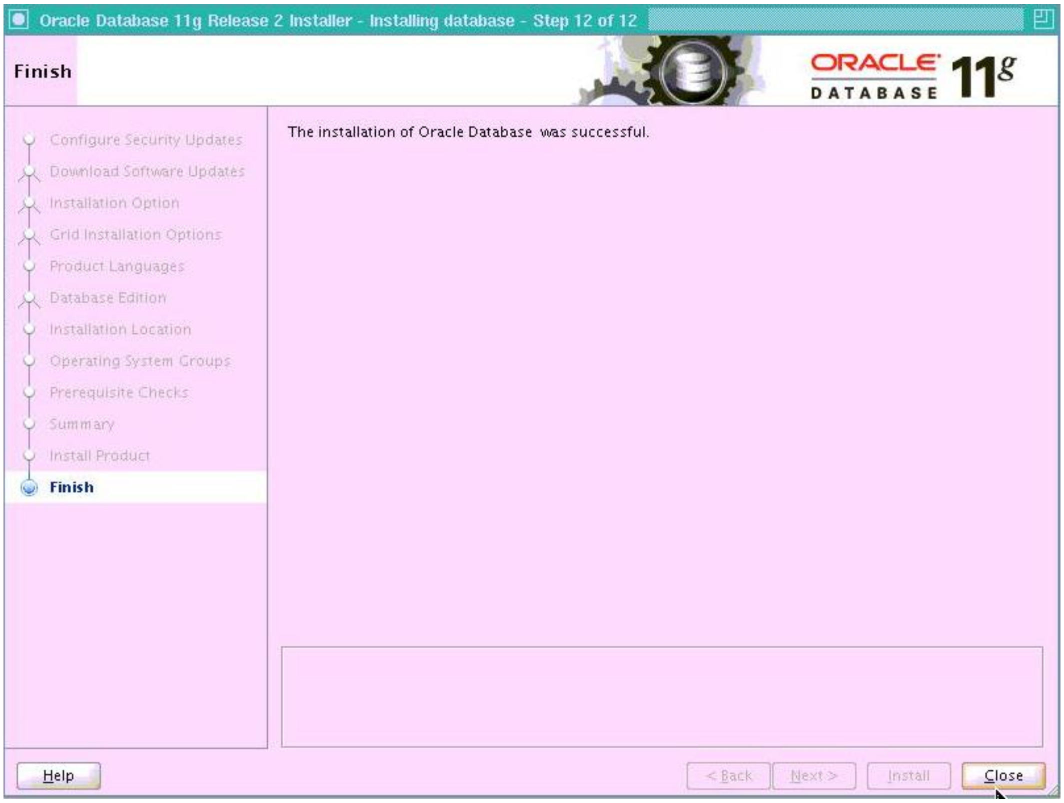

The installation is complete.

15. Close the Oracle Universal Installer.

5.3 Creating RAC Database on Cluster Nodes

Once the Oracle binary installation is completed, create your database by choosing to create NFS volumes on the storage controller and exporting it to be accessed by the NFS clients. In this scenario, an NFS client is the RAC node itself. During the database creation, all the volumes that are being used must be mapped with NFS v3 protocol. Direct NFS can be enabled after the database is created.

Description
Prerequisite description: Oracle 11*g* R1 (11.1) or later is required for using Oracle DNFS.

To create a RAC database on cluster nodes, refer to the following documentation:

http://docs.oracle.com/cd/E11882_01/install.112/e24660/dbcacrea.htm

Direct NFS can use up to four network paths defined for an NFS server. The Direct NFS client performs load balancing across all specified paths. If a specified path fails, Direct NFS reissues I/Os over any remaining paths.

5.4 Enabling DNFS

To configure DNFS, complete the following steps.

1. Direct NFS determines mount point settings to NFS storage devices based on the configurations in /etc/mnttab. Direct NFS looks for the mount point entries in the following order:

```
$ORACLE_HOME/dbs/oranfstab
/etc/oranfstab
/etc/mnttab
```

2. Shut down the database and the relevant ASM instance.
3. Execute the following `make` command to enable DNFS.

```
make -f $ORACLE_HOME/rdbms/lib/ins_rdbms.mk dnfs_on
```

Or

4. Perform the following steps to enable DNFS.
   a. Change to LD_LIBRARY_PATH.

```
cd $ORACLE_HOME/lib
```

   b. Rename the original file libodm11.so to libodm11.so.ORIG.

```
mv libodm11.so libodm11.so.ORIG
```

   c. Create a symbolic link from the standard ODM library to point to the NFS ODM library. This replaces the standard ODM driver with the ODM NFS library.

```
ln -s libnfsodm11.so libodm11.so
```

   d. Once the setup is complete, list the file to confirm that it is pointing to the right file.

```
$Host1 > ls -l libodm11.so
lrwxrwxrwx 1 oracle oinstall 14 Sep 5 09:02 libodm11.so -> libnfsodm11.so
```

5. Start the ASM instances on all the cluster nodes and then the databases.

6 Workload and Database Configuration

Benchmark Factory is a database performance and code scalability testing tool that simulates users and transactions on the database and replays a simulated workload in nonproduction environments. A benchmark is a performance test of hardware and/or software on a system under testing. Benchmark Factory provides the option of using industry-standard benchmarks during the load testing process. Benchmarks measure system peak performance when performing typical operations.

The two scenarios tested in this experiment are:

1. Deployment of four-node Oracle Database 11*g* R2 GRID infrastructures with RAC option in a virtualized environment
2. Scaling and consolidation study of several two-node Oracle Database 11*g* R2 GRID infrastructures with RAC option in a virtualized environment

Figure 6) Guest VM in each blade server.

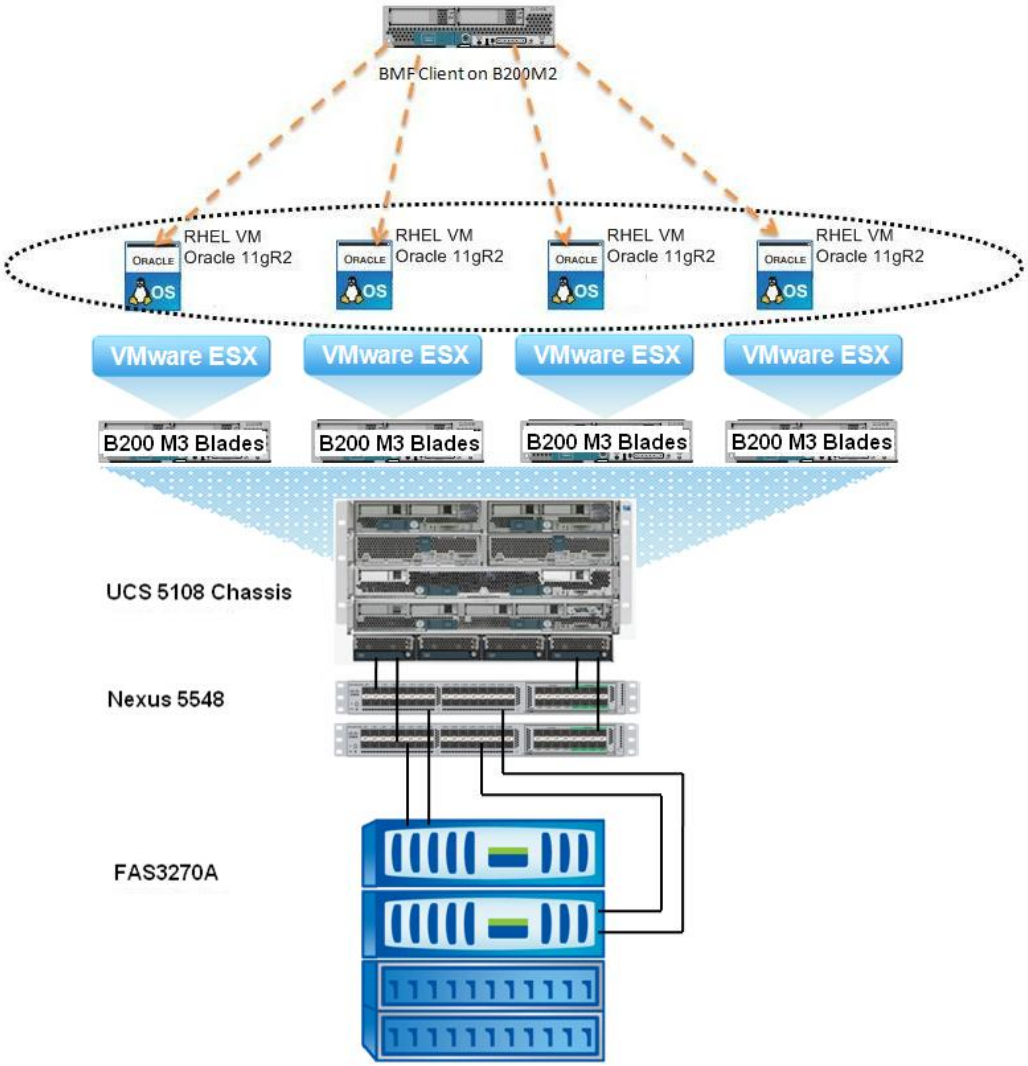

Oracle Database binary and Oracle Database 11*g* R2 GRID infrastructure with RAC option are installed in the local storage of each guest VM. However, these binaries can be installed in shared NFS storage. The database's data file, redo log files, CRS, and voting disks are stored in the NetApp storage system and are accessed using NFS.

The Oracle Direct NFS client is configured after installation of the database binary and database creation.

Figure 6 illustrates the configuration of the directory structure created in each guest VM to store the database as well as database binaries.

6.1 Scaling and Consolidation Study of Several Two-Node RAC VMs

This experiment concentrates on the scaling and consolidation of several two-node Oracle Database 11*g* R2s with RAC option on FlexPod using VM-FEX and Cisco VICs. Two fully loaded B200 M3s (Cisco blade servers) were used.

Three types of consolidation scenarios for the Oracle 11*g* R2 Database with RAC option can be implemented. However, only scenario 1, in which one guest VM was created on each blade to set up the two-node Oracle Database 11*g* R2 with RAC option, was used in this experiment. After loading the database on the newly set up Oracle clustered database, successive VMs were added on each of the ESXi servers to set up the individual two-node Oracle Database 11*g* R2s with RAC option.

Figure 7 illustrates scenario 1, in which the Oracle 11*g* R2 Database with RAC consolidation is configured with one RAC instance on each blade.

Figure 7) Scenario 1: Guest VM in each blade server with various Oracle Cluster configurations.

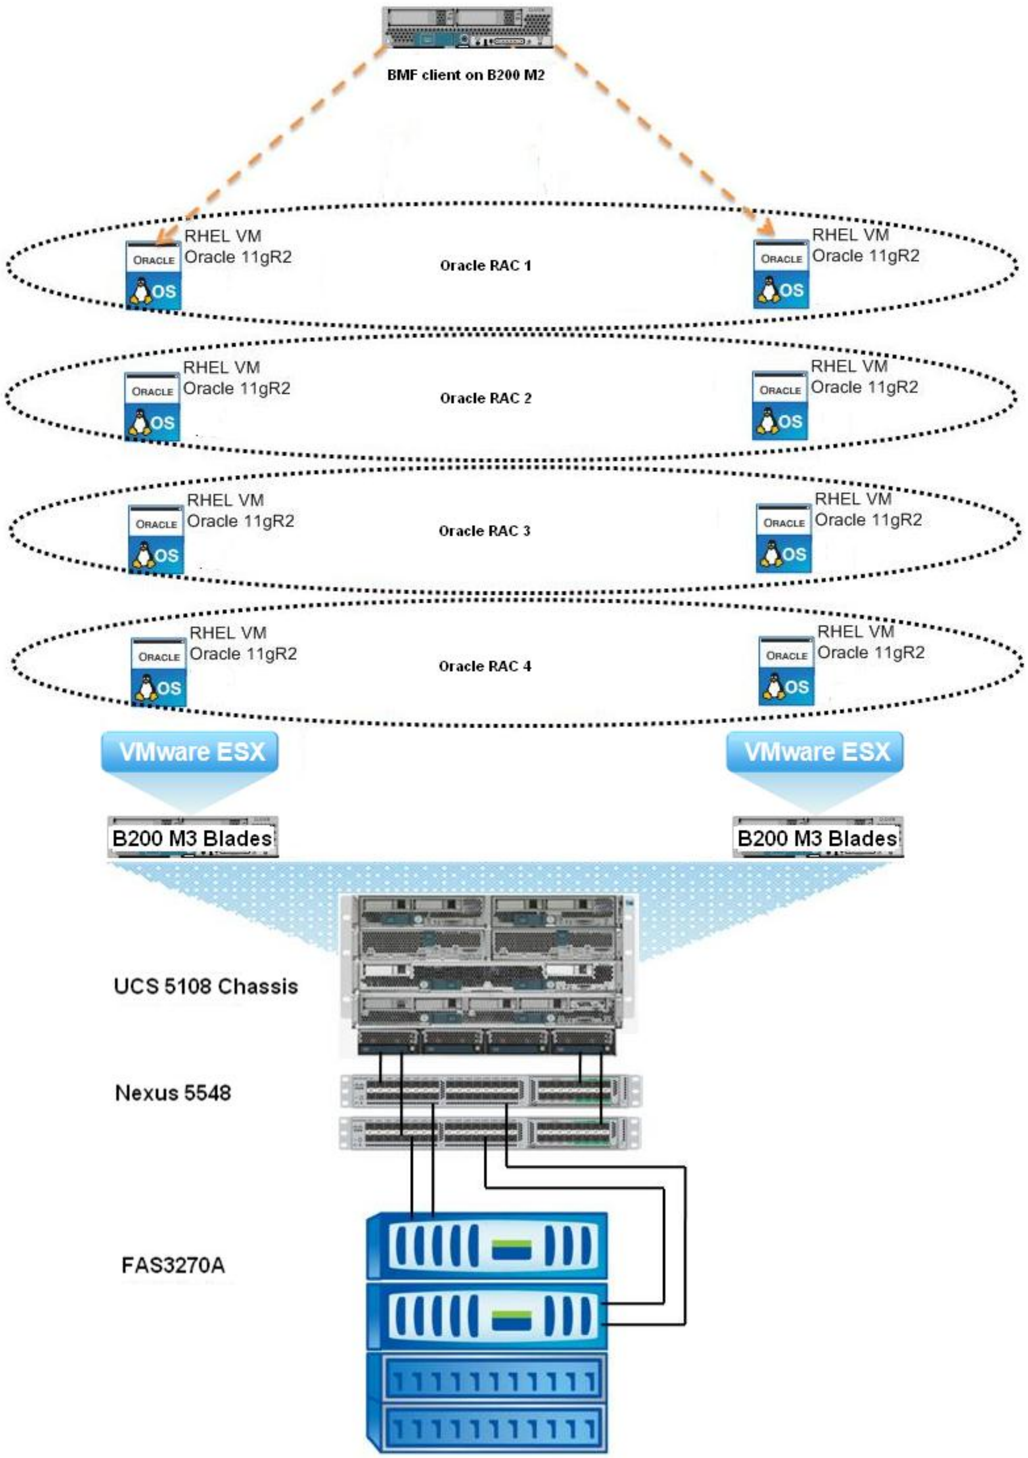

Figure 8 illustrates scenario 2, Oracle 11*g* R2 Database with a large RAC consolidation with a guest VM configuration.

Figure 8) Scenario 2: Oracle 11*g* R2 Database with RAC consolidation.

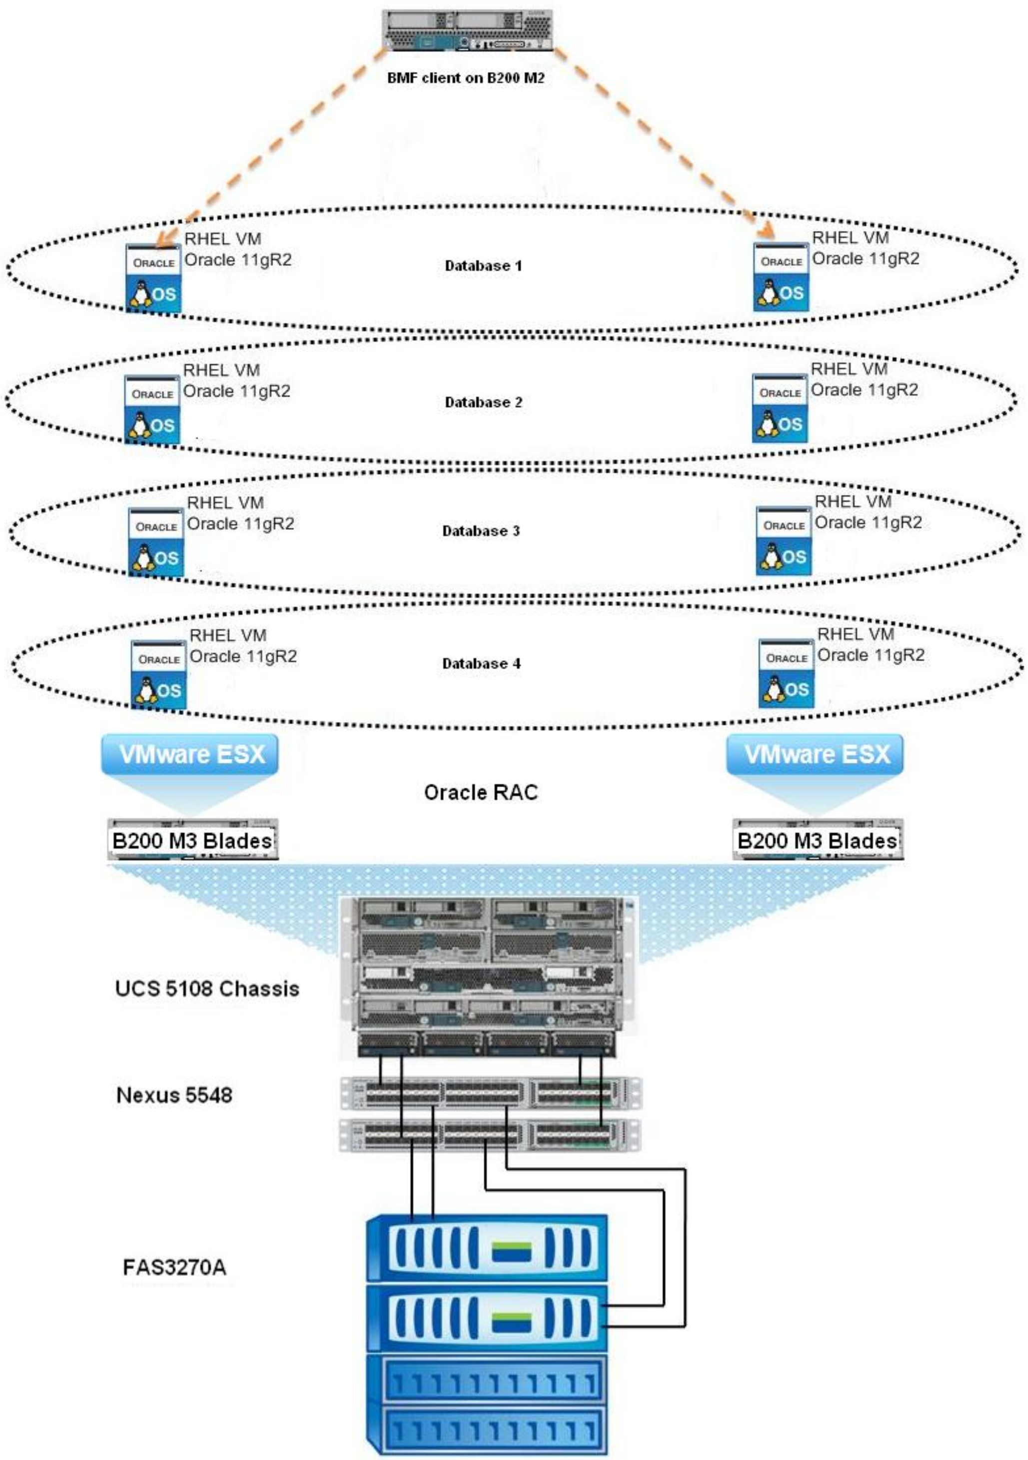

Scenario 3, Oracle 11g R2 Database with multiple databases in a two-node RAC configuration is illustrated in Figure 9.

Figure 9) Scenario 3: Several Oracle 11*g* R2 Databases with RAC consolidation.

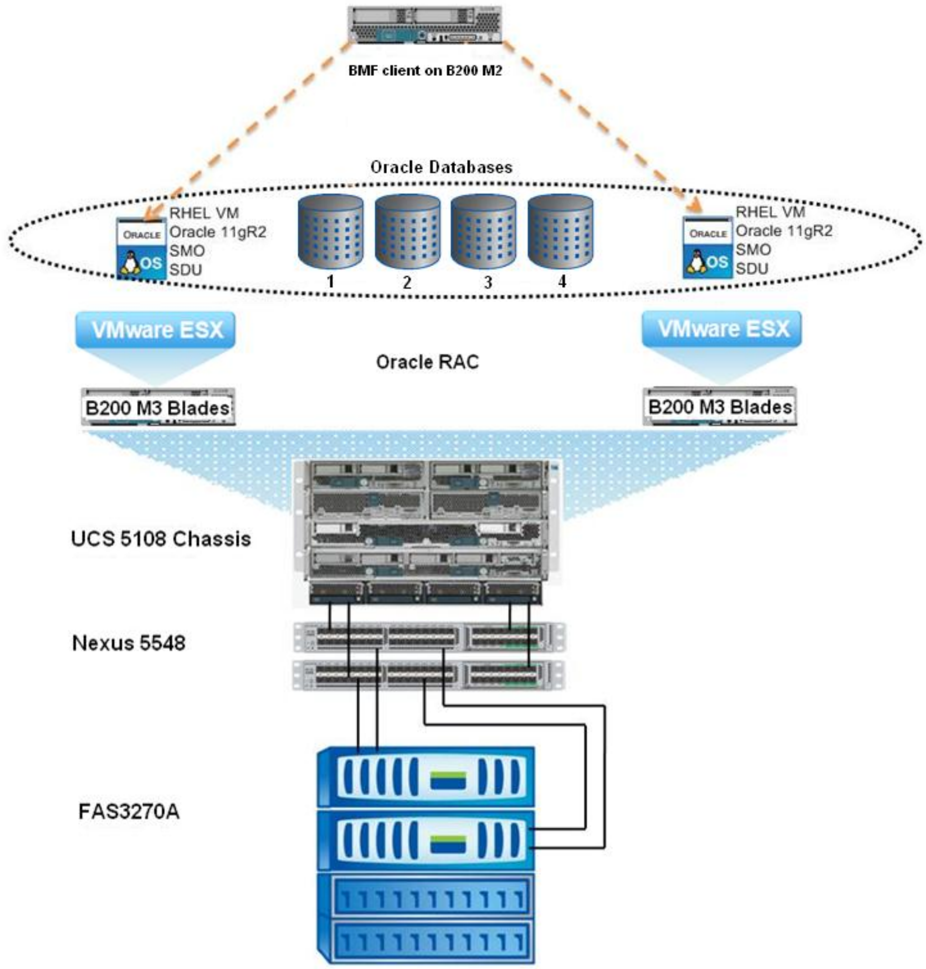

The other two scenarios can be implemented as per requirements. Oracle Database binary and Oracle GRID infrastructure are installed in the local storage of each guest VM; however, these binaries can be installed in shared NFS storage. Database data files, redo log files, CRS, and voting disks are stored in NetApp storage and accessed through NFS.

Table 3 describes the configuration of the directory structure created in each guest VM to store databases as well as binaries.

Table 3) Guest VM local directory configuration in each Oracle Cluster.

Guest VM Directory	Details	Description
/u01/app/oracle	Locally created	Database binary
/u01/app/11.2.0/	Locally created	GRID infrastructure
/ocrvote	Mounted with NFS volume	OCR file and voting disk
/data_A	Mounted with NFS volume from Controller A	Data files and copy of control file
/data_B	Mounted with NFS volume from Controller B	Data files and copy of control file
/log_A	Mounted with NFS volume from Controller A	Redo log files and copy of control file
/log_B	Mounted with NFS volume from Controller B	Redo log files and copy of control file

Note: A 500GB-sized database was created for each Oracle Cluster using Benchmark Factory for databases.

Table 4 describes the storage configuration required for each Oracle Cluster to store 500GB-sized databases.

Table 4) Storage configuration for database.

Aggregate Name	Storage Controller	Volume Name	Size	Description
Oracle_Aggr	Controller A	/vol/Data_Vol	250GB	Data files and copy of control file
Oracle_Aggr	Controller A	/vol/Log_Vol	100GB	Redo logs and copy of control file
Oracle_Aggr	Controller A	/vol/ocrvote	20GB	OCR file and voting disk
Oracle_Aggr_B	Controller B	/vol/Data_Vol_B	250GB	Data files and copy of control file
Oracle_Aggr_B	Controller B	/vol/Log_Vol_B	100GB	Redo logs and copy of control file

7 Conclusion

The Oracle RAC built on FlexPod with VMware is a balanced configuration that integrates servers, network, and storage in a single solution. The predictable sizing of this solution provides a useful approach to quickly determine ways to overcome performance issues and mitigate risks associated with new deployments. The appropriate sizing for the solution was achieved by the hard work of Cisco and NetApp professionals and the design is capable of handling concurrent workloads from different dimensions.

The pretested and validated solution accelerates deployment time and drastically lowers costs. With access to database clones in a mixed infrastructure environment, development teams can engage in parallel processes and reduce development and test cycle durations. This solution offers clear advantages in the areas of VM cloning and mass deployment, database cloning and data replication, and storage efficiency. It significantly improves savings in time and cycles associated with the process of cloning VMs and databases, and it delivers significant savings in storage consumption when compared to traditional storage and cloning methods. The benefits of using the NetApp Oracle application development and test solution are faster application deployments, increased revenues, and a competitive advantage over other competitors in the industry.

NetApp support for several protocols provides customers with the desired level of flexibility that is necessary for enabling a dynamic IT environment. Customers rely on NetApp in the implementation and planning stages to support numerous protocols and to provide data on how these protocols might perform in their environment in various situations (for example, moving a volume, failover).

The NetApp for Oracle Database solution delivers true business value to organizations regardless of their storage requirements.

8 Bill of Materials

Table 5 provides details of the components used in this document.

Table 5) Component descriptions.

Description	Part Number
UCS 5108 Blade Server Chassis	N20-C6508
UCS 2208XP I/O Module (8 External, 32 Internal 10Gb Ports)	UCS-IOM-2208XP
UCS B200 M3 Blade Server	UCS-FI-6248UP

Description	Part Number
Dual Intel® Xeon® E5-2690 CPUs (2.7 GHz and 8 cores) 256GB RAM (DDR3 1600 MHz)	
UCS 6248UP 1RU Fabric Int/No PSU/32 UP/ 12p LIC	UCS-FI-E16UP
UCS 6200 16-port expansion module/16 UP/ 8p LIC	FAS3270A
FAS3270 single-enclosure HA (single 3U chassis)	X1139A-R6
Dual-port 10GbE Unified Target Adapter with fiber	DS4243-1511-24S-QS-R54
Disk shelf with 600GB SAS drives, 15k RPM, 4 PSU, 2 IOM3 modules	SW-T7C_NFS-C
NFS software license	N5K-C5548UP-FA
Cisco Nexus 5548up	N5548P-SSK9
Cisco Nexus 5548up Storage Protocols Services License	SFP-10G-SR
10GBASE-SR SFP module	

Table 6 provides details of the software used in this document.

Table 6) Software details.

Platform	Software Type
UCS 6248UP	Management
UCS 6248UP	OS
Cisco Nexus 5548UP	OS
NetApp FAS3270	OS
Blade servers	OS
GRID and database	Oracle

References

The following documents were used as references in this guide.

- Cisco UCS
  http://www.cisco.com/en/US/netsol/ns944/index.html
- VMware vSphere
  http://www.vmware.com/products/vsphere/overview.html
- Oracle Databases on VMware Best Practices Guide
  http://www.vmware.com/files/pdf/partners/oracle/Oracle_Databases_on VMware_-Best_Practices_Guide.pdf
- TR-3749: NetApp Storage Best Practices for VMware vSphere
  http://media.netapp.com/documents/tr-3749.pdf
- NetApp Storage Systems
  http://www.netapp.com/us/products/storage-systems/
- TR-3932: Red Hat Enterprise Linux Protocol Performance Comparison with Oracle Database 11*g* Release 2
  http://media.netapp.com/documents/tr-3932.pdf
- TR-3961: Oracle Database 11*g* Release 2 Performance Using Data ONTAP 8.1 Operating in Cluster-Mode
  http://media.netapp.com/documents/tr-3961.pdf
- Data ONTAP 8.1 Software Setup Guide
  https://library.netapp.com/ecm/ecm_get_file/ECMP1114493

- Data ONTAP 8.1 System Administration Guide
  https://library.netapp.com/ecm/ecm_get_file/ECMP1113925
- Data ONTAP 8.1 Data Protection Guide
  https://library.netapp.com/ecm/ecm_get_file/ECMP1113990
- Data ONTAP 8.1 Network and File Access Management Guide
  https://library.netapp.com/ecm/ecm_get_file/ECMP1113296
- Cisco Nexus
  http://www.cisco.com/en/US/products/ps9441/Products_Sub_Category_Home.html
- Cisco Validated Design—FlexPod for VMware
  http://www.cisco.com/en/US/docs/solutions/Enterprise/Data_Center/Virtualization/flexpod_vmware.html
- Cisco White Paper: Oracle RAC Built on FlexPod with VMware
  http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/whitepaper_C11-717661.html
- Cisco Nexus 5000 Series NX-OS Software Configuration Guide
  http://www.cisco.com/en/US/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide.html
- TR-3298: RAID-DP: NetApp Implementation of Double Parity RAID for Data Protection
  http://www.netapp.com/us/library/technical-reports/tr-3298.html

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

NetApp®
www.netapp.com