

Deploying and Migrating Oracle RAC to Oracle VM with Dell EMC SC Series Storage

Dell EMC Storage Engineering
July 2017

Revisions

Date	Description
May 2016	Initial release
July 2017	Minor updates on document references

Acknowledgements

Author: Henry Wong, Product Specialist

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2016 - 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [7/18/2017] [Best Practices] [2022-A-RA-O]

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

Revisions.....	2
Acknowledgements.....	2
1 Introduction.....	5
1.1 Audience.....	5
1.2 Overview.....	6
2 Oracle VM component specifications.....	8
2.1 Oracle VM Servers	8
2.2 Oracle VM Manager	12
2.3 Networks.....	18
2.4 Storage	20
3 Oracle VM configuration.....	31
3.1 Create a server pool	31
3.2 Create repositories and present to a server pool	32
3.3 Create networks.....	33
3.4 Create virtual machines for Oracle RAC	34
3.5 Create anti-affinity group	36
3.6 Install the guest OS	36
3.7 Configure VM network adapter.....	37
3.8 Configure iSCSI initiator	39
3.9 Configure multipathing and udev rules	39
3.10 Add VMs to DSM	40
4 Install Oracle Grid Infrastructure and DB software.....	41
4.1 Perform Oracle DB preinstallation tasks	41
4.2 Install the Grid Infrastructure and RDBMS software	42
4.3 Create ASM disk groups.....	44
4.4 Create a RAC database	44
5 Migrate Oracle RAC database from physical servers to VMs.....	45
5.1 Steps to migrate a RAC database	46
1 Terminology.....	50
1 Configuration files.....	52
1.1 Volume group definition for boot volume on Oracle VM Manager server	52
1.2 Volume group definition for vgoracle	54

1.3	Multipath configuration on Oracle VM Servers.....	55
1.4	Multipath configuration on Oracle VM Manager server.....	58
2	Technical support and resources	60
2.1	Related documentation.....	60

1 Introduction

Oracle® VM is a robust virtualization solution designed to manage and run virtual machines (VMs) with a variety of applications and workloads. Many enterprise customers are looking to virtualization to consolidate multiple enterprise software and hardware on fewer servers. With a single Oracle VM Server, customers can run multiple VMs with different operating systems and software, improving resource utilization while decreasing software and utility costs.

Oracle Real Application Cluster (RAC) is a cluster database that is typically deployed on physical servers. Oracle has contributed significant resources to continuously advance Oracle VM with new and improved features, better scalability, manageability, and has streamlined the user interface of the management tool. Therefore, it has become more common to deploy Oracle RAC on Oracle VM.

This document presents an Oracle VM configuration running Oracle RAC database and covers the use of Dell EMC™ SC Series arrays in such a virtualized environment with Oracle RAC. The advanced features of SC Series arrays provide further enhancement to the applications running on Oracle VM and Oracle RAC. These features include but are not limited to volume thin provisioning, automatic storage tiering, and volume snapshots.

- This document covers the topics for Oracle VM version 3.x for the Linux® x86 platform only. It does not cover the installation and administration of the Oracle software. The official Oracle VM documentation library contains detailed information on typical installation and configuration. Refer to [Oracle Virtualization Product Documentation Libraries](#) for more details.

It is also recommended to review the [Dell SC Series Storage Best Practices for Oracle VM](#) document which covers many Dell Storage best practices and guidelines. Some of this information is also presented here in this document.

Note: The information contained within this document is intended to be general recommendations and may not be adaptable to all configurations. The actual configuration may vary based upon individual or business needs.

1.1 Audience

This document is intended for system administrators, database administrators, and other information technology professionals who are interested in setting up an Oracle VM environment for Oracle RAC.

This document assumes the reader has understanding of the following:

- Operation and configuration of SC Series arrays
- Oracle Linux or Red Hat® Enterprise Linux operating systems
- Oracle VM software installation and virtualization concepts
- Oracle Grid Infrastructure and RAC software installation and administration

1.2 Overview

This section describes the architecture of the Oracle VM for x86 solution, which is depicted in Figure 1. The specification of each component is listed in section 2.

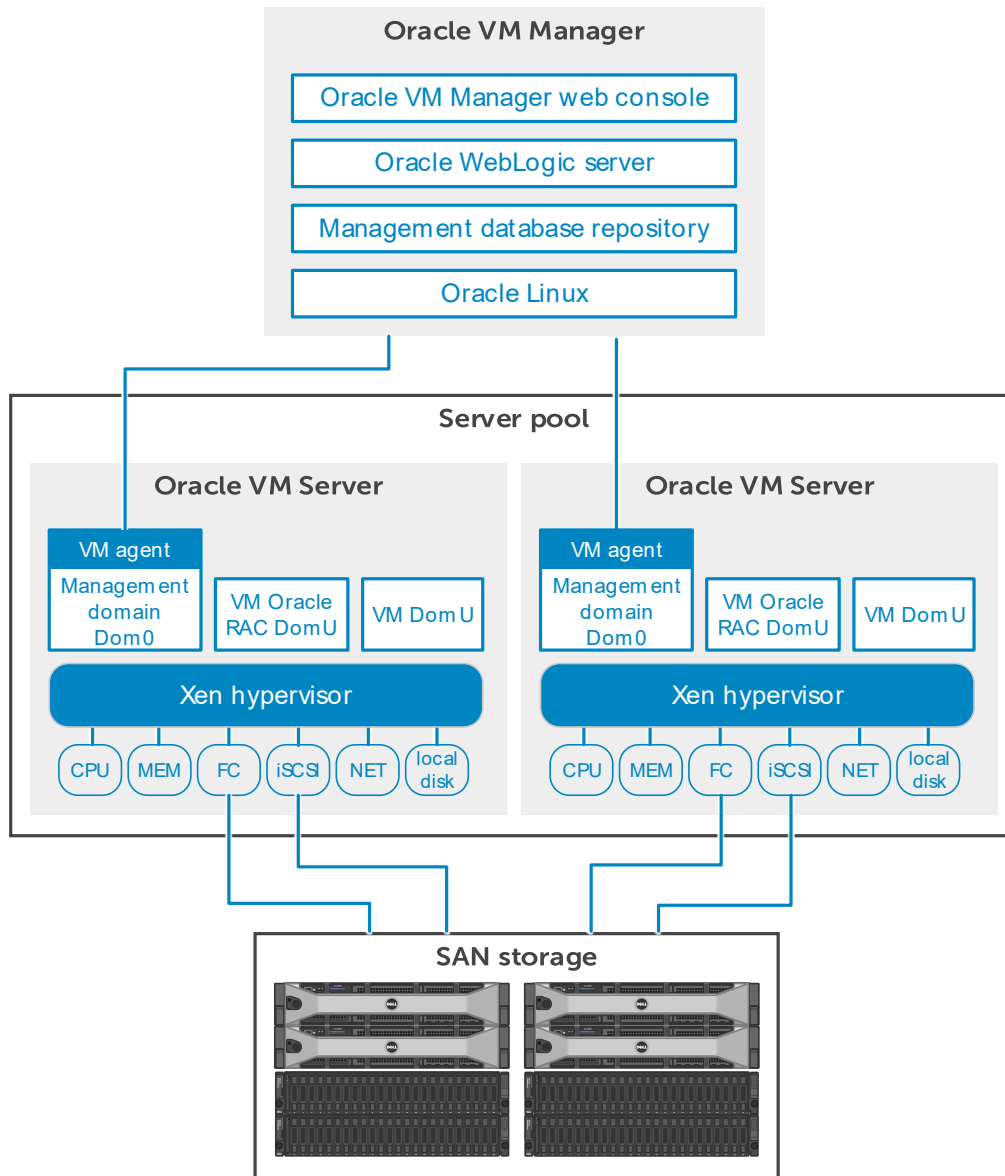


Figure 1 Architecture of Oracle VM

Oracle VM Server for x86: This runs on a 32-bit or 64-bit x86 architecture hardware. The hypervisor is based on open-source Xen technology which provides an abstraction layer for all hardware operations. The Oracle VM agent runs on each Oracle VM Server inside the management domain, Dom0, which has direct access to the hardware such as the processor, memory, local drives, and host bus adapters (HBAs). It is responsible to manage and monitor all the hardware resources and virtual machine guests. Oracle VM Servers are usually clustered together and form server pools.

Oracle VM Server for x86 supports a wide range of guest operating systems. They include Oracle Linux, Red Hat® Enterprise Linux, Oracle Solaris, CentOS, Microsoft® Windows® and Windows Server®. For up-to-date information on supported guest operating systems, consult the *Oracle VM Release Notes* on the [Oracle documentation](#) site.

Oracle VM Manager: Oracle VM Manager is a web-based management tool that runs on Oracle Linux on a separate server. It is built on the Oracle WebLogic application and requires a database repository backend to store configuration information. It provides a friendly and centralized interface to manage all aspects of Oracle VM Servers including the storage area network (SAN) storage. In addition to the web interface, it also has a command-line interface that allows managing the environment through scripting.

Oracle VM guests: These are VMs running one of the supported operating systems such as Oracle Linux, Red Hat Enterprise Linux, or Microsoft Windows. Multiple VMs can run on an Oracle VM Server and can migrate to other Oracle VM Servers in the same cluster server pool where the same hardware resources are available. For a two-node Oracle RAC, two VMs are set up on different Oracle VM Servers for load balancing and high availability.

Shared storage: Third-party NFS, FC, and iSCSI SAN storage are supported through the storage-connect framework. It provides a shared space for holding files and data used by virtual machines and the cluster database for the server pool.

Both FC and iSCSI shared storage are presented in this document for different uses to Oracle VM Servers and VMs.

Refer to the [Oracle VM Concepts Guide](#) for further details.

2 Oracle VM component specifications

This section describes each Oracle VM component specification including servers, network, storage, and software.

2.1 Oracle VM Servers

The Oracle VM cluster consists of two Dell PowerEdge™ R620 servers. These servers form a server pool in which the virtual machines run. Oracle VM Server 3.3.3 is installed on the servers.

2.1.1 Processor and memory

When choosing the model of servers, consider the processor, memory, and network requirements of all virtual machines plus the control domain dom0. The physical server should have enough capacity to run all the virtual machines combined at the same time.

In addition to the virtual machines, additional memory is required for the control domain dom0. Oracle specifies the following algorithm for dom0 memory allocation:

$$\text{dom0_mem} = 502 + \text{int}(0.0205 * \text{Physical Memory})$$

The memory allocation might need to be increased when a large number of FC or iSCSI LUNs are used on the Oracle VM Servers. Oracle recommends 1 MB per LUN in addition to the memory required to run dom0. The dom0_mem can be changed in the /boot/grub/grub.conf file on each Oracle VM Server:

```
kernel /xen.gz console=com1,vga com1=57600,8n1 dom0_mem=max:2112M allowsuperpage
```

Table 1 summarizes the server hardware and BIOS settings. Enable the virtualization technology in the BIOS if the hardware supports the feature.

Table 1 Oracle VM Server configuration used

Component	Description
Model	PowerEdge R620
Processors	2 x 8-core Intel® Xeon® Processor E5-2670 @ 2.60GHz
Memory	32GB
Swap	32GB
OS	Oracle VM Server release 3.3.3
Hardware architecture	X86_64
Boot volume	100GB FC SAN volume
BIOS System Profile	Performance
BIOS Logical Processor	Enabled
BIOS QPI Speed	Maximum Data Rate

Component	Description
BIOS Virtualization Technology	Enabled
BIOS Logical Processor Idling	Disabled
BIOS Memory Operating Mode	Optimizer Mode

2.1.2 Boot volume

Oracle VM Server supports both local disk and SAN volume for the boot volume. A multipathed SAN boot volume is configured with the partition layout shown in Table 2. Review the layout and sizes and adjust them according to the actual need and available space.

Note: Per Oracle VM Server requirements, a SAN boot volume must be a multipathed volume.

Table 2 Oracle VM Server boot volume partitions

Device	Mount point	Size
/dev/mapper/36000d31000fc000000000000000053p1	/boot	1GB
/dev/mapper/36000d31000fc0000000000000000053p2	/	25GB
/dev/mapper/36000d31000fc0000000000000000053p3	swap	32GB

2.1.3 Network adapter

Oracle VM Servers require at least one network for server management, cluster heartbeat, and other VM-related communications. Multiple networks can be configured to isolate different types of traffic. In the physical server, six network interfaces are present in dom0. They are configured into different networks and presented to the VMs.

Each server has a quad-port Intel® Ethernet 10G 4P X540 adapter and a dual-port Intel Ethernet server adapter X520-2.

Table 3 shows the network interfaces and their purposes. See the section 2.3 for additional network information.

Table 3 Network interfaces on Oracle VM Servers

Network interface	Adapters	Network
bond0	<ul style="list-style-type: none"> • 1GB minimum • 2 x 10GB is preferred • Use of Linux bonding to provide port redundancy and automatic failover • Static IP assignment 	Server management, Cluster Heartbeat, Live Migration
bond1	<ul style="list-style-type: none"> • 1GB minimum • 2 x 10GB is preferred • Use of Linux bonding to provide port redundancy and automatic failover • Static IP assignment 	Public virtual machine communication
bond2	<ul style="list-style-type: none"> • 1GB minimum • 2 x 10GB is preferred • Use of Linux bonding to provide port redundancy and automatic failover • Static IP assignment • The Clusterware network should be isolated from the public server management and iSCSI networks 	Oracle Clusterware communication between virtual machines
eth2	<ul style="list-style-type: none"> • 10GB • iSCSI network 1 should be on a different VLAN than iSCSI network 2 • Static IP assignment 	iSCSI network 1
eth3	<ul style="list-style-type: none"> • 10GB • iSCSI network 1 should be on a different VLAN than iSCSI network 2 • The two iSCSI networks provide SAN storage redundant paths • Static IP assignment 	iSCSI network 2

Due to the limited number of adapter ports on the physical server, there is only one port assigned to a bond interface in the example. Ideally, two ports are desired to form a bond interface to provide high availability and failover.

Note: If using bond interfaces, set them up in the Oracle VM manager interface. Do not modify the `/etc/sysconfig/network-scripts/ifup*` files directly because it will cause Oracle VM manager to be out of sync.

To set up the bond interface, navigate to the server under the **Server Pool**, and change the **Perspective** to **Bond Ports**. See Figure 1.

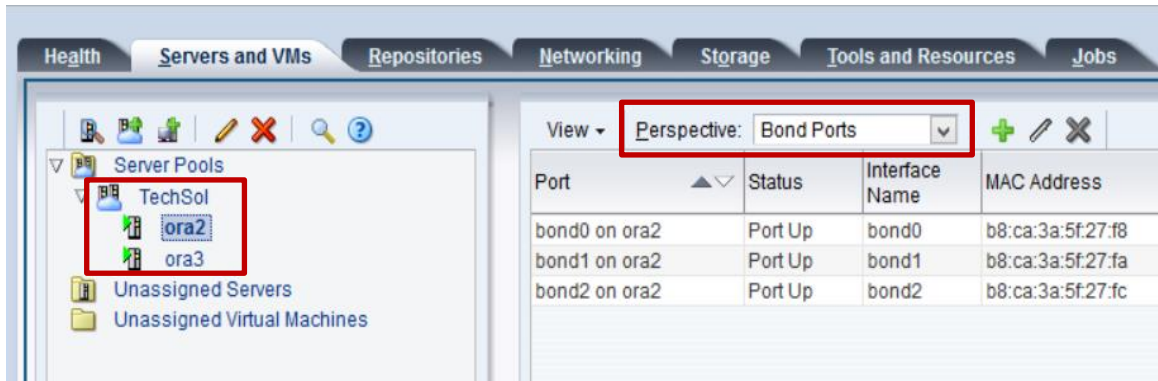


Figure 1 Bond ports set up in Oracle VM

2.1.4 Fibre Channel adapter

Each Oracle VM Server has a QLogic® QLE2562 PCI-E dual-channel 8Gb FC HBA. The dual-FC port configuration provides path redundancy to the SAN.

2.1.4.1 QLogic BIOS settings

The following BIOS settings (Table 4) are set by the Fast!UTIL BIOS utility. Enter the utility during the system boot process. Make sure the settings are adjusted on all adapter ports.

Table 4 QLogic BIOS settings

Parameter	Value
Connection Options (point to point only)	1
Login Retry Count	30
Port Down Retry Count	5
Link Down Timeout	60
Execution Throttle	256

2.1.4.2 Host bus adapter queue depth and module settings

The default queue depth value of the QLogic adapter is typically set to 32. While the optimal queue depth will vary depending on a number of factors, a value of 64 or 128 is common and works well in most cases.

This value was modified on the Oracle VM Servers by updating the module file, `/etc/modprobe.d/ql2xxx.conf`, with the following:

```
options qla2xxx ql2xmaxqdepth=128
```

For more information on setting HBA parameters on the Oracle VM Servers, reference the documents, [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 6x Best Practices](#) and [Dell SC Series Storage Best Practices for Oracle VM](#). The Oracle VM Server is based on Oracle Enterprise Linux 6 which is very similar to Red Hat Enterprise Linux 6.

2.1.5 Multipathing

For Oracle VM Servers for Linux x86, Device-Mapper Multipath (DM-Multipath) packages are installed and enabled by default. This is a requirement of Oracle VM Manager allowing it to discover the SAN disks.

Each SAN volume, including the boot volume, has four access paths on the FC SAN. The default multipath configuration is implemented without modification on the Oracle VM Servers. The **device** section in `/etc/multipath.conf` file should already include the following specific settings for **COMPELNT**.

```
defaults {
    user_friendly_names no
    getuid_callout "/lib/udev/scsi_id --whitelisted --replace-whitespace --device=/dev/%n"
}

device {
    vendor                "COMPELNT"
    product               "Compellent *"
    path_grouping_policy  multibus
    path_checker          tur
    failback              immediate
    rr_min_io             1024
    no_path_retry         10
}
```

A full listing of the `/etc/multipath.conf` is included in appendix 1.3.

Note: It is an Oracle VM requirement to set the `user_friendly_names` parameter to **no** in the `/etc/multipath.conf` file.

2.2 Oracle VM Manager

The Oracle VM Manager is a web-based management console that manages and monitors the Oracle VM Servers and virtual machines running on them. A separate server is typically deployed to run the applications.

The Oracle VM Manager supports Red Hat Enterprise Linux or Oracle Enterprise Linux, version 5U5 and version 6 and above. Dell EMC recommends running the latest version of Oracle Linux 6 or Red Hat Enterprise Linux 6 if possible and apply the standard Linux best practices. For details, refer to the document, [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 6x Best Practices](#).

In this document, a PowerEdge R620 server was deployed with Oracle Linux 6.7.

2.2.1 Processor and memory

The minimum server requirements include one 64-bit x86 processor and 8 GB of memory. Table 5 lists the server configuration used for this document.

Table 5 Oracle VM Manager server configuration

Component	Description
Model	PowerEdge R620
Processors	2 x 8-core Intel Xeon Processor E5-2670 @ 2.60GHz
Memory	32GB
Swap	32GB
OS	Red Hat Enterprise Linux 6.7
BIOS system profile	Performance
Hardware Architecture	x86_64
Boot volume	100G FC SAN volume
BIOS system profile	Performance
BIOS Logical Processor	Enabled
BIOS QPI Speed	Maximum data rate
BIOS Virtualization Technology	Enabled
BIOS Logical Processor Idling	Disabled
BIOS Memory Operating Mode	Optimizer Mode

2.2.2 Boot volume

The server was configured to boot from an SC Series array on the FC SAN. The boot volume must have the **Map volume using LUN 0** option enabled during the mapping in Dell EMC Storage Manager (DSM) as shown in Figure 2. A major advantage to use a SAN boot volume is that it can leverage the storage snapshot to protect the operating system. More information on this can be found in the document, [Dell SC Series Storage Best Practices for Oracle VM](#).

The following volume and server will be mapped

Volume ovs-ora3-boot

Server ora2

Select LUN

☐ Use next available LUN
LUN to use when mapping to Volume

☒ Use the next available LUN if specified LUN is unavailable

☒ Map volume using LUN 0 (this is usually reserved for boot volumes)

Restrict Mapping Paths

☒ Allow the Storage Center to automatically determine the Controller to activate the Volume on

☒ Map to All Available Server Ports

Configure Multipathing

Maximum number of paths per Server

Configure Volume Use

☐ The volume should be presented as read-only to the server

Figure 2 Map volume advanced option for boot volume

Logical volume manager (LVM) was configured on the boot volume. Disk partitions can also be used but LVM offers better flexibility in space management. Table 6 shows the volume group layout used for this document. The detailed volume group definition for vg00 is listed in appendix 1.1.

Table 6 Oracle VM Manager server boot volume partitions

Volume Group	Logical volume	Physical volume	Mount point	Size
vg00	lv_root	/dev/mapper/mpathap2	/	50G
vg00	lv_home		/home	18G
vg00	lv_swap		SWAP	32G
Disk partition	N/A	/dev/mapper/mpathap1	/boot	477M

vg00 consisted of a multipathed physical volume `/etc/mapper/mpathap2` with four access paths. The number of paths available varies depending on the HBA configuration on the server and the SC Series array.

```
mpatha (36000d31000fcaf0000000000000000054) dm-4 COMPELNT,Compellent Vol
size=100G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
    |- 6:0:2:0   sdc 8:32   active ready running
    |- 7:0:0:0   sdk 8:160  active ready running
    |- 6:0:1:0   sdb 8:16   active ready running
    `-- 7:0:1:0   sdl 8:176  active ready running
```

2.2.3 File system for applications

Oracle recommends installing the Oracle VM Manager applications in the `/u01` file system. The minimum size for `/u01` is 5.5GB. Table 7 shows the `/u01` file system configuration. The detailed volume group definition for `vgoracle` is listed in appendix 1.2.

Table 7 Oracle VM Manager application file system

Volume group	Logical volume	Physical volume	Mount point	File system type	Size
vgoracle	lv_u01	/dev/mapper/mpathb	/u01	ext4	20G

The `vgoracle` volume group consists of one multipathed SC volume with four access paths. The benefit of putting the file system on LVM is that it can be easily extended in the future without service interruption when the needs arise.

```
mpathb (36000d310000065000000000000000201f) dm-1 COMPELNT,Compellent Vol
size=20G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
    |- 6:0:5:3   sdg 8:96   active ready running
    |- 7:0:3:3   sdo 8:224  active ready running
    |- 6:0:6:3   sdi 8:128  active ready running
    `-- 7:0:4:3   sdq 65:0   active ready running
```

2.2.4 Network adapter

The server had a quad-port Intel Ethernet 10G 4P X540 adapter and a network bond interface was set up for communicating to the Oracle VM Servers on the public management network. Port bonding provides port-level redundancy for the server. Table 8 shows the network adapter configuration.

Table 8 Network interfaces on Oracle VM Manager server

Network interface	Adapters	Network
bond0	<ul style="list-style-type: none">• 2 x 1GB minimum (10GB is preferred)• Use of Linux bonding to provide port redundancy and automatic failover• Static IP assignment	<ul style="list-style-type: none">• Management network for accessing and managing Oracle VM Servers• Oracle VM Web GUI access

2.2.5 Fibre Channel adapters

The same QLogic QLE2562 adapter was installed in the Oracle VM Manager server. The settings of this adapter can be found in section 2.1.4.

2.2.6 Multipathing

DM-Multipath packages were installed and enabled to provide path failover and load-balancing management. Four access paths were available for each SAN volume on the FC SAN. In the `/etc/multipath.conf` file, set **user_friendly_names** parameter to **yes**. Under the **COMPELNT** device section, set **path_checker** to **tur**, which is default, for best performance. A complete listing of the `/etc/multipath.conf` file can be found in appendix 1.4.

```
defaults {
    user_friendly_names yes
}

device {
    vendor                "COMPELNT"
    product               "Compellent *"
    path_grouping_policy  multibus
    path_checker          tur
    failback              immediate
    rr_min_io             1024
    no_path_retry         10
}
```


2.2.7 Oracle VM Manager applications

Oracle VM Manager version 3.3.3 consists of the following applications:

- Oracle VM Manager web GUI
- Oracle WebLogic Server 12.1.2.0
- MySQL 5.6.24 enterprise commercial advanced database for the management repository

Note: Oracle VM Manager version 3.2 and above provides and installs its own copy of MySQL software. It is fully supported in a production environment. If there is a preinstalled copy of MySQL on the operating system, it needs to be removed to avoid conflicts and confusions.

2.2.7.1 Prerequisite and optional rpms

In addition to the Oracle Linux 6.7 **Database Server** software bundle, the following rpm packages were installed for this document:

Prerequisite RPMs: These rpms are required by Oracle VM Manager and must be installed before the Oracle VM Manager installation. Although Oracle did not specify the versions explicitly, Dell EMC recommends installing the latest version of the packages.

- `zip`
- `unzip`
- `libaio`
- `libaio-devel`
- `perl`

Optional RPMs: These additional rpm packages were not required by Oracle VM Manager but were useful for managing the environment.

- **System time update:** Allows updates and a synchronized system clock to time servers
 - `ntpdate`
 - `ntp`
- **SCSI disk utilities:** Provides useful utilities such as `rescan-scsi-bus.sh` to rescan all SCSI devices
 - `sg3_utils-libs`
 - `sg3_utils`
- **iSCSI software initiator:** Allows bringing iSCSI SAN storage to the system
 - `iscsi-initiator-utils`
- **Apache Web Server:** Provides a mean to import ISOs to Oracle VM Storage repository
 - `httpd`
 - `httpd-tools`

2.3 Networks

This section presents the network infrastructure for the Oracle VM cluster and Oracle RAC solution. The following networks were set up for different purposes:

Management network: This is used for the Oracle VM cluster heartbeat and other management functions. Oracle VM Manager must be able to communicate to the Oracle VM Servers on this network. The SC Series array also uses it for management access.

Public network: This is used for inter-VM communication and outside access to the public for all VMs and servers.

Oracle RAC cluster network: Required by Oracle RAC, this network is strictly for RAC cluster communication between the VMs. Traffic should be isolated from other networks.

iSCSI networks: This is used by Oracle RAC in the VMs. The cluster crs/voting disk and the database volumes resided on SC Series array and were accessed through the iSCSI networks. More information on iSCSI SAN is in section 2.4.1.

Table 9 summarizes the network types and their corresponding interfaces across all components.

Table 9 Networks summary

Network	Oracle VM Servers	Oracle VM Manager	Oracle RAC VMs	SC Series array	Speed
Management network	bond0	bond0	N/A	N/A	1G
Public network for servers, VMs and Storage Center	bond1	bond0	eth0	Controller 1 MGMT Port Controller 2 MGMT Port	1G
Oracle RAC cluster network	bond2	N/A	eth1	N/A	1G
iSCSI network 1	eth2	N/A	eth2	Controller 1, Slot 3, Port 1 Controller 2, Slot 3, Port 1	10G
iSCSI network 2	eth3	N/A	eth3	Controller 1, Slot 3, Port 2 Controller 2, Slot 3, Port 2	10G

A Dell EMC Networking 5548 switch was used for connecting the servers and SC Series array. The management network, public network, and Oracle RAC cluster network shared the same switch and VLANs were created to separate the traffic.

Note: Although only a single switch was used in this document, it is recommended to use two switches to provide switch-level redundancy.

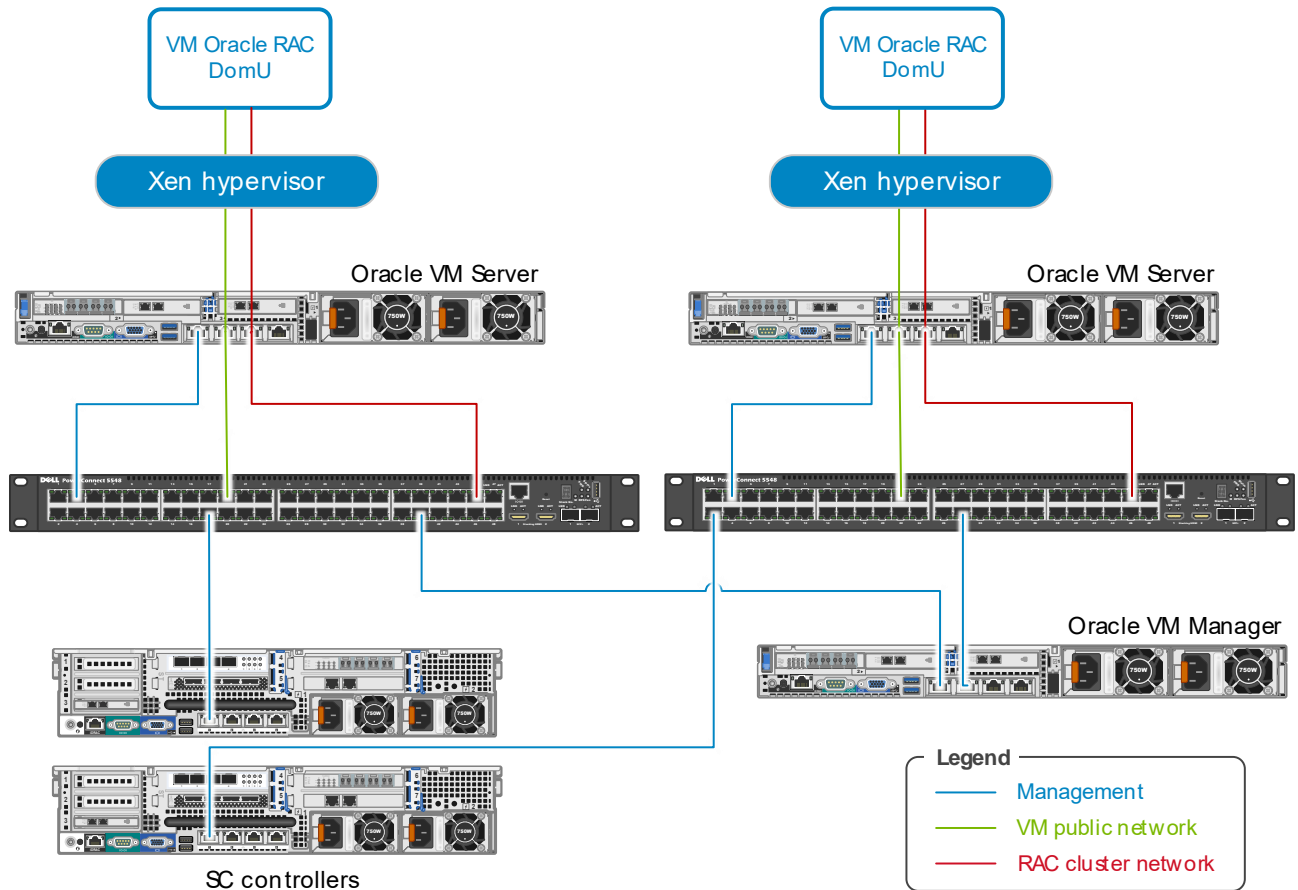


Figure 3 Networking diagram

Table 10 Dell EMC Networking 5548 configuration

Component	Description
Model	Dell EMC Networking 5548
OS	4.1.0.19
Jumbo Frames	Disabled
Flow Control	Enabled on switch ports
Spanning Tree Mode	RSTP on edge ports
VLANs	4000: Private cluster network for RAC 550: Management, public network

Support documentation on the Dell EMC Networking switch can be found on the [Dell Support](#) site.

2.4 Storage

Oracle VM supports both Fibre Channel and iSCSI SANs. To demonstrate the usage of both SANs, the Oracle VM storage repository was set up to reside on the FC SAN while the Oracle Linux VMs accessed the iSCSI SAN for the database volumes.

For more information on storage types, storage options, and connectivity, refer to the [Oracle virtualization product documentation libraries](#) and the document, [Dell SC Series Storage Best Practices for Oracle VM](#).

2.4.1 iSCSI SAN

A Dell EMC Networking S5000 network switch was configured to handle the iSCSI traffic between the PowerEdge servers and the Dell SC Series array. To isolate the iSCSI traffic from the public, two VLANs were established and assigned different IP subnets.

Note: Although only a single switch was used for this document, it is recommended to use two switches to provide switch-level redundancy.

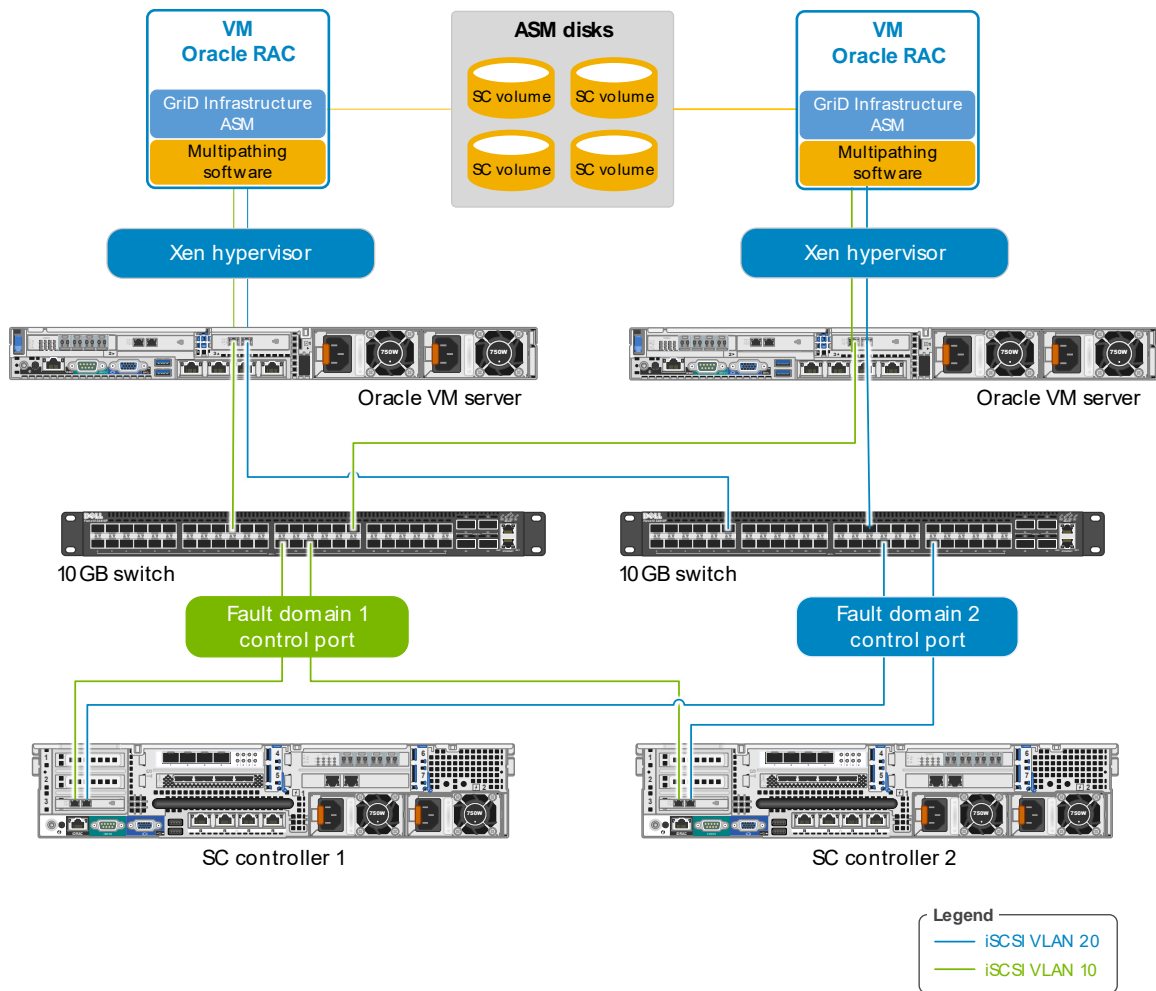


Figure 4 iSCSI SAN with dual networks, dual controllers, control ports, and virtual fault domains

The S5000 switch was configured based on the best practices in the document, [Dell Networking S5000 Switch Configuration Guide for SC Series iSCSI SANs](#). Enable bi-directional flow control on all server and switch ports that handle iSCSI traffic. Also, set spanning-tree mode and Data Center Bridging according to the best practices.

Table 11 10G networking switch specification

Component	Description
Model	Dell EMC Networking S5000
OS	9.8
Jumbo Capable	Yes
Module	4 x 12-port 10GE SFT+
Flow Control	Enabled on switch ports

Spanning Tree Mode	RSTP on edge ports
VLANs	10 – iSCSI network 1 20 – iSCSI network 2 550 – Management network
Data Center Bridging	Disabled

Additional support documents can be found on the Dell TechCenter page, [Dell Networking S5000](#), and on [Dell.com/support](#). For other switch models, see the page, [Switch Configuration Guides for PS Series or SC Series SANs](#).

Oracle VM Servers: Each Oracle VM Server is connected to both iSCSI VLANs with two 10G connections. Each Ethernet port must be enabled in Oracle VM Manager, but it is not necessary to assign IP addresses to the ports because they would not mount SC volumes directly in this scenario. Instead, Oracle VM Servers present the iSCSI networks to the VMs.

View ▾ Perspective: Ethernet Ports ▾						
Port	△ ▾	Status	Interface Name	MAC Address	Bond Name	Network
eth3 on ora2		Port Up	eth3	00:1b:21:d4:ef:fd		iscsi-vlan20
eth2 on ora2		Port Up	eth2	00:1b:21:d4:ef:fc		iscsi-vlan10

View ▾ Perspective: Ethernet Ports ▾						
Port	△ ▾	Status	Interface Name	MAC Address	Bond Name	Network
eth3 on ora3		Port Up	eth3	00:1b:21:71:b1:4d		iscsi-vlan20
eth2 on ora3		Port Up	eth2	00:1b:21:71:b1:4c		iscsi-vlan10

Figure 5 Ethernet ports information in Oracle VM Manager

Dell EMC recommends enabling virtual port mode on the SC Series array controllers to provide port-level and controller-level redundancy. Each controller is connected to both 10Gb iSCSI VLANs. Each VLAN is associated with an iSCSI fault domain which consists of ports from the same VLAN. Find more information on the fault domain in section 2.4.3.1.

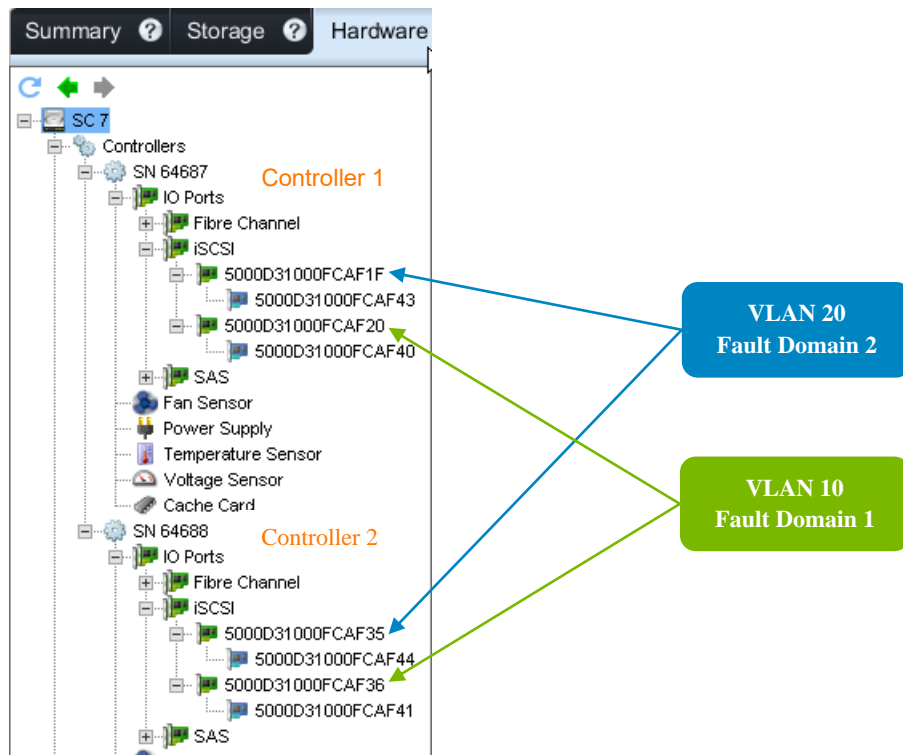


Figure 6 SC array iSCSI VLANs and fault domains

2.4.2 Fibre Channel SAN

The FC SAN provides storage for Oracle VM storage repositories. It consists of two fabrics with multiple Brocade® 6505 (24 ports 16Gbps) switches. Dell EMC recommends running the latest stable version of the firmware on the switch. N_Port ID Virtualization (NPIV) must be enabled on the ports. FC fault domains must be set up on the SC Series array for high availability. See section 2.4.3.3 for more information on FC fault domains.

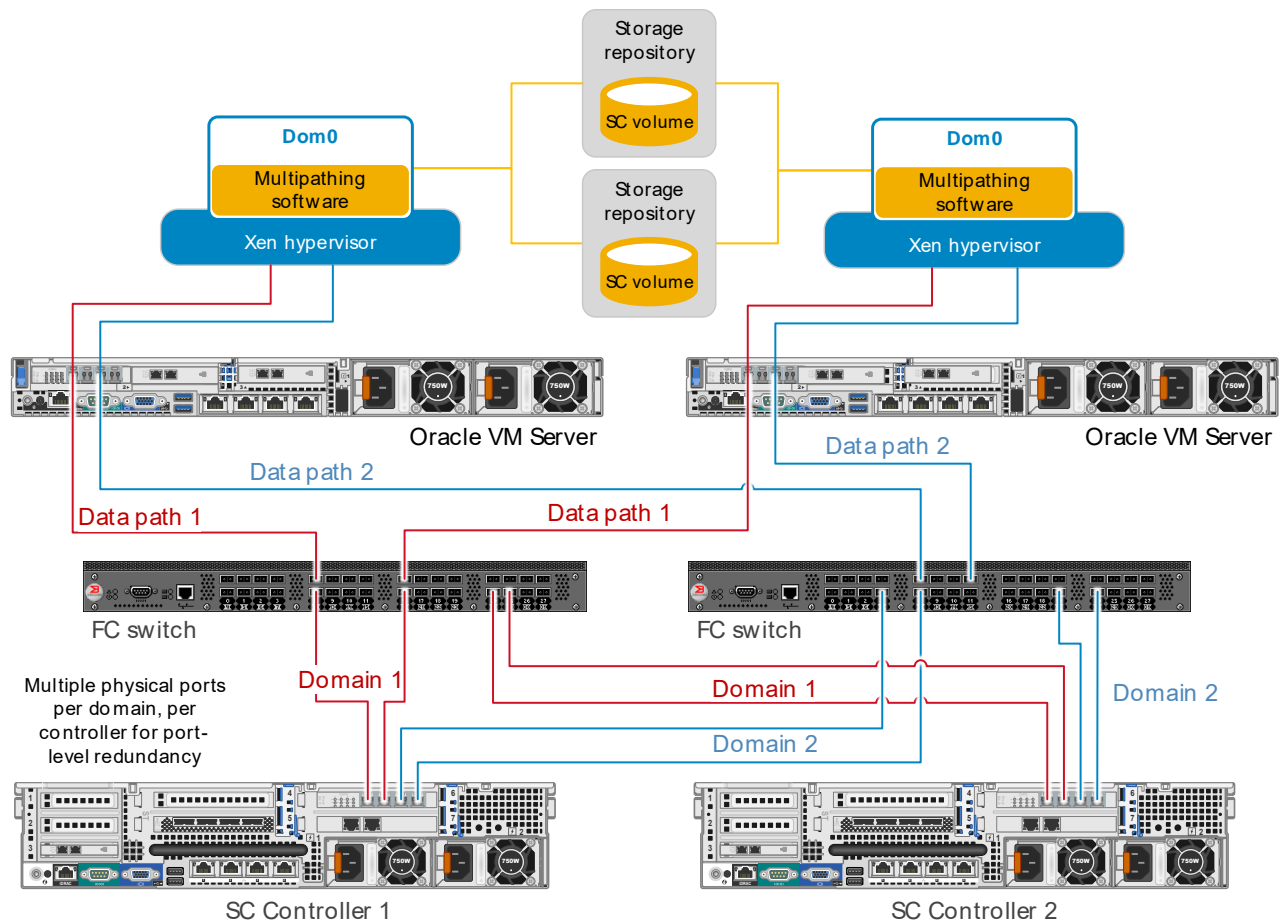


Figure 7 Fibre Channel SAN with dual fabrics, dual controllers, and virtual fault domains

2.4.2.1 Zoning

Fibre Channel zones are used to segment the fabric to restrict access. A zone contains paths between initiators (server HBAs) and targets (storage array front-end ports). When deploying an SC Series array on an FC SAN, follow specific zoning guidelines relating to SC Series arrays. For more information, refer to the [Dell SC Series Storage Best Practices for Oracle VM](#) on Dell TechCenter, and the Dell Storage Center System Manager Administrator's Guide available on the Knowledge Center at the [SC Series customer portal](#) (login required).

A summary of FC zoning guidelines is as follows:

- Use name zoning
- Create a zone with all physical WWNs on the SC Series array controllers (see Figure 8)
- Create a zone with all virtual WWNs on the SC Series array controllers (see Figure 8)
- Use two server zones for each Oracle VM Server and Oracle VM Manager server
 - Use single-initiator, multiple-target zones for server zones
 - Include one server FC port and the SC virtual port zone in each server zone

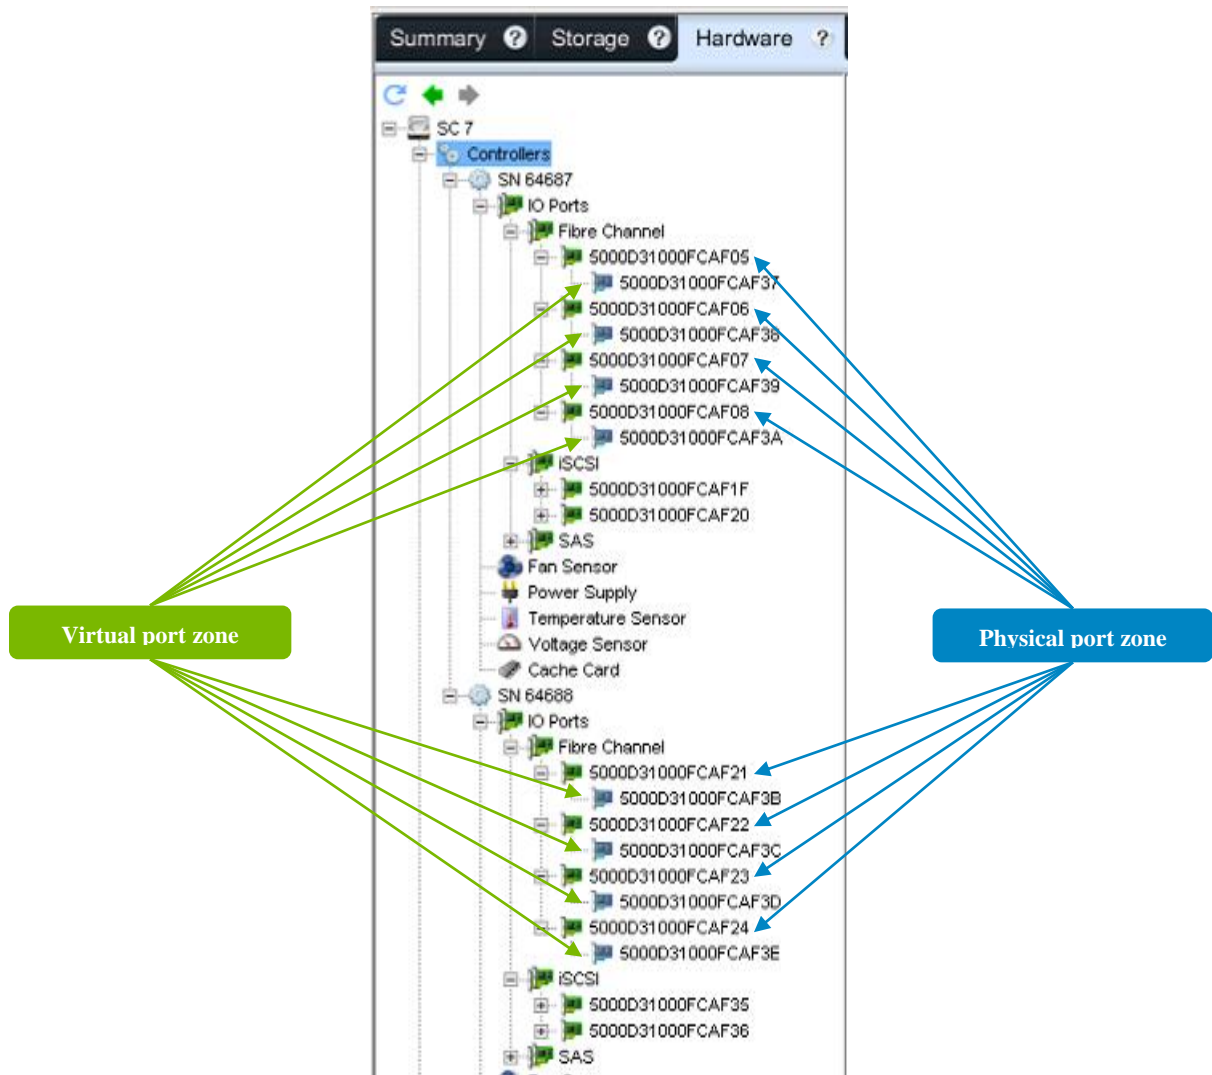


Figure 8 SC Series array FC physical and virtual ports

2.4.3 SC Series array

This section describes the configuration of the SC Series array deployed for this document.

A Dell SC8000 Series array was deployed to provide external SAN storage in an Oracle VM solution. It is a high-performance, space-efficient, and scalable storage solution that supports simultaneous iSCSI and FC SANs. When configuring the storage for an Oracle VM solution, consider its many advanced features to take advantage of all the benefits. For more details, refer to the document, [Dell SC Series Storage Best Practices for Oracle VM](#).

Table 12 lists the specifications for the SC8000 array.

Table 12 SC8000 array specification

Component	Description
Model	SC8000
Controllers	2
SCOS	6.7.5.5
Memory	56GB per controller
Front end ports	4 x 8 Gb FC ports per controller 4 x 10 Gb iSCSI ports per controller
Back-end ports	4 x 24 Gb SAS ports per controller
Port mode	Virtual Port Mode
Fault domain	2 FC fault domains 2 iSCSI fault domains
Disk enclosure	1 x SC220
Disks	Tier 1: 12 x 372GB write-intensive SSDs Tier 3: 12 x 930GB 7K drives
Power supply	Dual power supply per controller

2.4.3.1 Virtual ports and fault domains

Dell EMC recommends enabling virtual ports in the SC Series array. With virtual ports, all front-end I/O ports are active, thus increasing performance throughput, and allow automatic port-level failover to improve availability. One virtual port is created for every physical port. Any virtual ports can fail over to any physical ports within the same fault domain.

A fault domain groups virtual ports of the same type across both controllers. All port members in the same domain must connect to the same fabric or VLAN.

For more information on virtual ports and fault domains, consult the *Dell Storage Center Storage Manager Administrator's Guide*.

2.4.3.2 iSCSI fault domain

One fault domain is created for each iSCSI VLAN and a control port is associated with each domain. Static IPs are assigned to each physical port and control port, but only the control port IPs are used in the iSCSI configuration on the servers and VMs.

Table 13 iSCSI fault domains and control ports

Fault domain	Control Port IP	Corresponding VLAN
Fault Domain 1	10.10.7.10	VLAN 10
Fault Domain 2	10.20.7.10	VLAN 20

iSCSI

iSCSI Transport Mode Virtual Port

Fault Domains Front End Ports

Name	Controller	Status	Slot	Slot Port	Fault Domain	IPv4 Address	Subnet Mask	Gateway IPv4 Address	iSCSI Name	Purpose
5000D31000FCAF20	SN 64687	Up	3	1	Fault Domain 1	10.10.7.11	255.255.0.0	0.0.0.0	iqn.2002-03.com.compellent:5000d31000fcaf20	Front End
5000D31000FCAF36	SN 64688	Up	3	1	Fault Domain 1	10.10.7.12	255.255.0.0	0.0.0.0	iqn.2002-03.com.compellent:5000d31000fcaf36	Front End
5000D31000FCAF1F	SN 64687	Up	3	2	Fault Domain 2	10.20.7.11	255.255.0.0	0.0.0.0	iqn.2002-03.com.compellent:5000d31000fcaf1f	Front End
5000D31000FCAF35	SN 64688	Up	3	2	Fault Domain 2	10.20.7.12	255.255.0.0	0.0.0.0	iqn.2002-03.com.compellent:5000d31000fcaf35	Front End

Figure 9 iSCSI fault domains port members

2.4.3.3 FC fault domain

For each FC fabric, a separate fault domain must be created in DSM. FC ports on both controllers connected to the same fabric are assigned to the same domain.

Fibre Channel

Fibre Channel Transport Mode Virtual Port

Fault Domains Front End Ports

Name	Controller	Status	Connected	Slot	Slot Port	Fault Domain	Purpose	Description
5000D31000FCAF05	SN 64687	Up	Yes	6	1	Domain 1	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF06	SN 64687	Up	Yes	6	2	Domain 1	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF21	SN 64688	Up	Yes	6	1	Domain 1	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF22	SN 64688	Up	Yes	6	2	Domain 1	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF07	SN 64687	Up	Yes	6	3	Domain 2	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF08	SN 64687	Up	Yes	6	4	Domain 2	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF23	SN 64688	Up	Yes	6	3	Domain 2	Front End	QLogic QLE2564 8G Fibre Channel Adapter
5000D31000FCAF24	SN 64688	Up	Yes	6	4	Domain 2	Front End	QLogic QLE2564 8G Fibre Channel Adapter

Figure 10 FC fault domain port members

2.4.4 Oracle VM storage repositories and server pool file system on SC Series array

Oracle VM stores various types of data and files in a storage repository. A minimum of one repository must be created. Oracle VM limits a single disk per storage repository but multiple repositories are allowed. Oracle

recommends a minimal of 10 GB for a storage repository and enough space should be allocated for virtual machines, templates, DVD/CD image (ISO) files, and other virtual machine resources.

For each clustered server pool, Oracle VM requires an OCFS2-based server pool file system to store the cluster information and to facilitate the cluster heartbeats. Store the server pool file system on a separate small volume. A SC volume is created for this server pool file system.

Table 14 Repository and server pool file system information

Repository	Content	Size	Access frequency
TechsolRepo	Virtual machine configuration files Virtual disks Virtual machine templates Virtual machine assembly files	1TB	Frequent
TechsolRepo2	DVD/CD image files (ISO)	200GB	Rarely
Serverpool file system	OCFS2 file system Cluster information of a server pool Served as a cluster heartbeat device	15GB	Frequent

An SC volume appears to the Oracle VM Servers like a regular disk, but underneath it is a virtualized volume comprised of blocks from multiple physical disks, disk types, storage tiers, and RAID levels. The Data Progression feature continuously evaluates and adjusts the data placement to archive optimum balance between performance and cost. The attributes of an SC volume can be modified to meet different application requirements.

Table 15 shows the SC volume attributes for each storage repository.

Table 15 SC volume attributes

Volume attributes	TechSolRepo volume	TechSolRepo2 volume	ServerPool volume
Name	TechSolRepoVol	TechSolRepo2Vol	srvpool_filesystem
Storage profile ¹	Recommended (All Tiers)	Low Priority (Tier 3)	Recommended (All Tiers)
Snapshot profile ²	Daily	Daily	Daily
Read/write cache	Enabled/disabled	Enabled/disabled	Enabled/disabled
Compression ³	Enabled	Enabled	Enabled

¹The SC storage profile applied to the SC volume determines where the blocks reside and how they migrate between the storage tiers. The Recommended (All Tiers) profile is applied to the TechSolRepo and ServerPool volumes. This profile makes use of all disk types and storage tiers (both SSDs and spinning drives). The Low Priority (Tier 3) profile applied to the TechSolRepo2 volume confines the data (ISO image files) to the storage tier that consists of slower, high-capacity spinning disks. Additional predefined storage profiles are available or custom profiles can be defined by the user.

²A Daily snapshot profile is applied to the SC volumes which creates a snapshot once a day on each volume. Data Progression also manages frozen pages by a snapshot and moves them to different storage type and tiers based on the storage profile.

³*Enabling compression can further reduce storage utilization. Not all data is eligible for compression. Only frozen pages by a snapshot or inaccessible pages due to new data written over it are eligible for compression. When SSDs are installed in a SC Series array, Dell EMC recommends disabling write cache on all volumes using SSDs to maximize performance for most applications.*

All SC volumes are thin provisioned — the actual storage is not consumed until data is written to them. The volume can also be expanded dynamically using DSM. Oracle VM supports online volume expansion for the storage repository without service interruption.

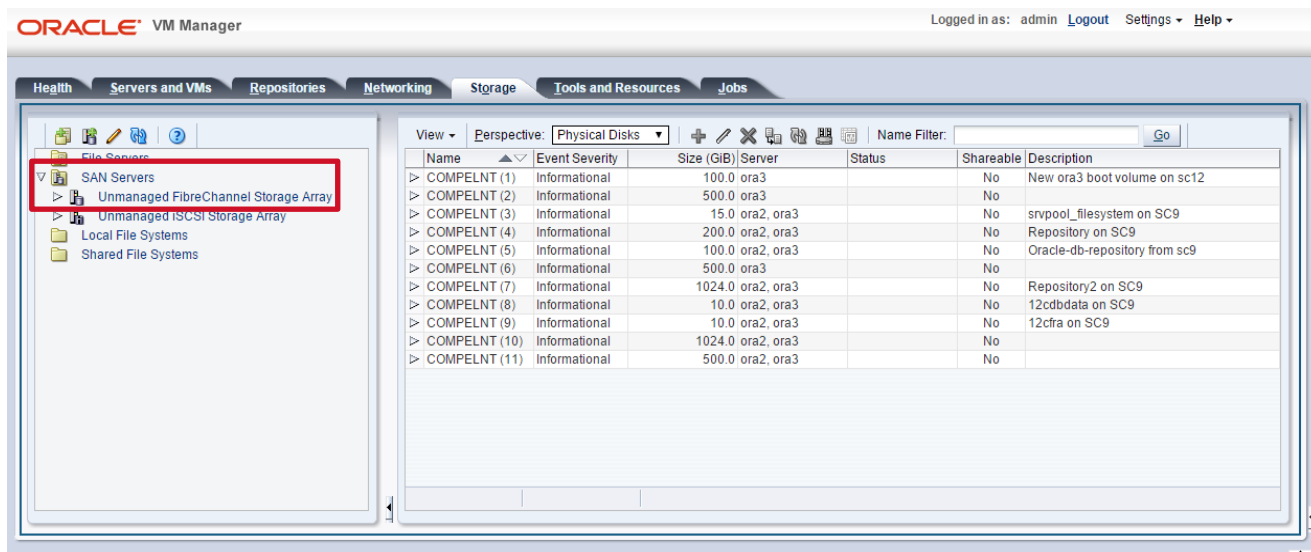
Automated storage tiering, thin-provisioning, Data Progression, and compression are transparent to Oracle VM. No configuration within Oracle VM is necessary to take advantage of these advanced features. For more information on these features, consult the *Dell Storage Center Storage Manager Administrator's Guide* and [Dell SC Series Storage Best Practices for Oracle VM](#).

Note: To maximize flexibility, performance, and cost-savings benefits, install multiple types of drives in an SC Series array and choose storage profiles that move data across all storage tiers, such as **Recommended (All Tiers)** or **Flash Optimized with Progression**.

2.4.4.1 Mapping SC volumes in Oracle VM Manager

This section describes the process to make the SC volumes available in Oracle VM Manager.

1. After mapping the volumes to the Oracle VM Servers in DSM, open the Oracle VM Manager web console to scan for the new disks.
2. Navigate to the **Storage** tab > **SAN Servers** folder. The SC volumes appear under the **Unmanaged FibreChannel Storage Array** or **Unmanaged iSCSI Storage Array**, depending on which transport protocol is used.



3. Add Oracle VM Servers to the **SAN Server** admin server list. Right-click the SAN server and edit the **Admin Servers** list.
4. Rescan and update storage on SAN servers. For example, for FC volumes, right-click **Unmanaged FibreChannel Storage Array** and select **Refresh**. Once the refresh is done, the new volumes show up under the **Physical Disks** perspective.

The newest additional disk appears at the bottom of the list. The name of the disk is automatically assigned with the vendor ID and an index number, though it can be changed afterward.

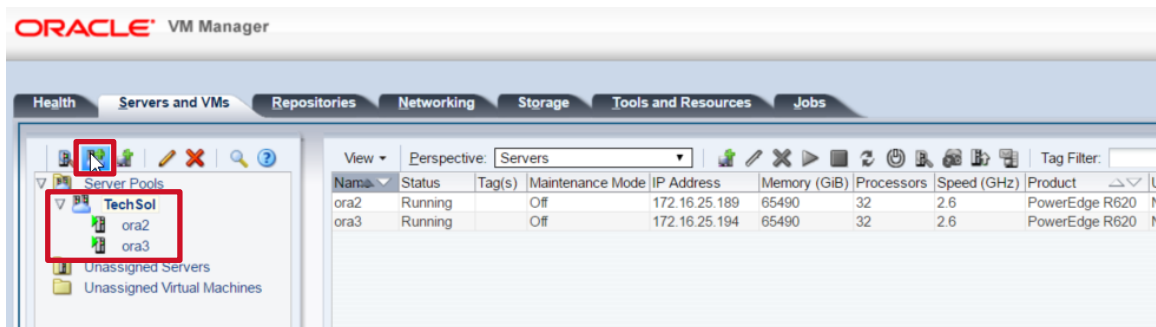
3 Oracle VM configuration

Before creating virtual machines, the following components must be configured in Oracle VM:

- A clustered server pool of two servers
- Storage repositories
- Networking

3.1 Create a server pool

To create a clustered server pool, under the **Servers and VMs** tab in the Oracle VM Manager web console, click the **create server pool** icon.

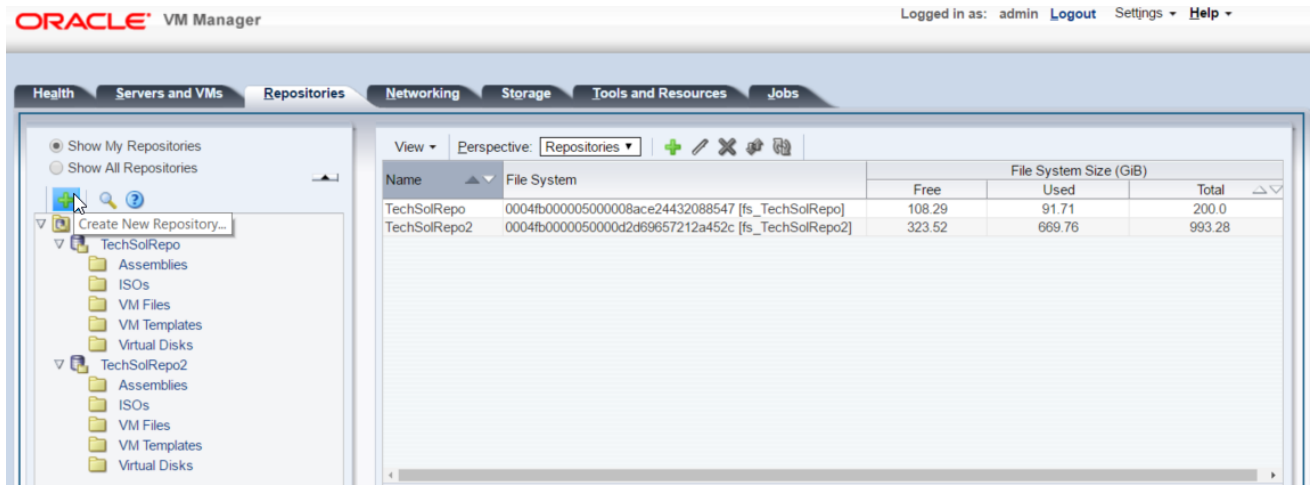


The server pool configuration used for this document is listed in the following table:

Parameter	Value
Name	TechSol
Virtual IP Address for the Pool	172.16.25.249
Assigned servers	ora2, ora3
Secure VM Migrate	Unchecked
Clustered Server Pool	Checked
VM Start Policy	Best Server
Timeout for Cluster	120
Storage for Server Pool	Physical Disk
Storage Location	Unmanaged FibreChannel Storage Array COMPELNT (3) (SC volume created in section 2.4.4)

3.2 Create repositories and present to a server pool

As discussed in section 2.4.4, two repositories are needed in this scenario. To create the repositories in Oracle VM Manager Web console, navigate to the **Repositories** tab and click the **+** icon in the left pane.

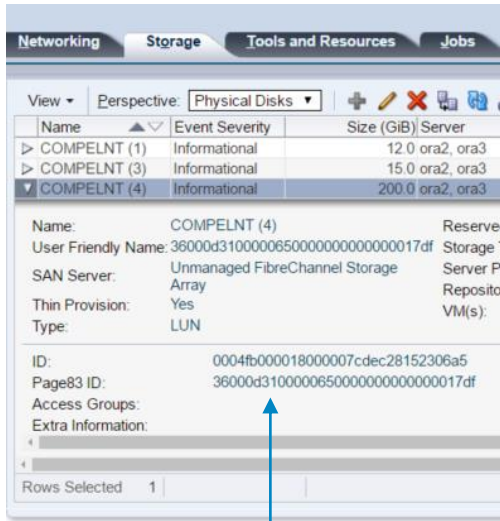


The values in the following table are used to create the repositories:

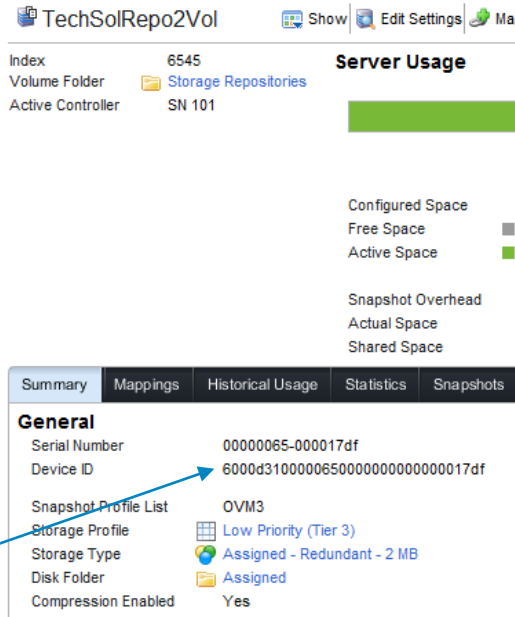
Parameter	TechSolRep	TechSolRep2
Repository Name	TechSolRep	TechSolRep2
Repository Location	Physical Disk	Physical Disk
Physical Disk	Unmanaged FibreChannel Storage Array1 COMPELNT (4) ² (SC volume created in section 2.4.4)	Unmanaged FibreChannel Storage Array COMPELNT (7) (SC volume created in section 2.4.4)
Server Pool	TechSol	TechSol

¹Oracle VM classifies SC Series array as Unmanaged FibreChannel Storage Array or Unmanaged SCSI Storage Array shows up in Oracle VM because the array cannot be managed directly from Oracle VM. The management of the array must be performed in DSM.

²To help identifying the SC volume, use the device page83 ID information to correlate with the SC volume device id in DSM. Page83 ID can be found under the Oracle VM Storage tab, Physical Disks perspective.



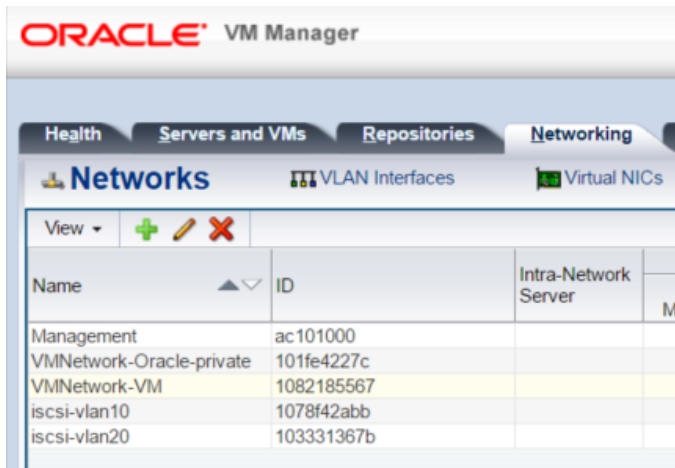
Correlate Page83 ID in Oracle VM with SC volume device id correlate



3.3 Create networks

Configure the **Ethernet Ports** and **Bond** interfaces on each Oracle VM Servers under the **Servers and VMs** tab, **Server Pools**, [Server Name]. Then, the new **Network** can be created. Port and network information was discussed in section 2.3.

To set up new networks, navigate to the **Networking** tab and click the **+** icon. The following networks were created based on Table 9 in section 2.3. The network has **Network Channels** selected for each network.



Network name	Ports	Network Channels	Description
Management ¹	bond0 on ora2, ora3	Server Management Cluster Heartbeat Live Migrate	Oracle VM management network
VMNetwork-Oracle-private	bond2 on ora2, ora3	Virtual Machine	Oracle RAC cluster network
VMNetwork-VM	bond1 on ora2, ora3	Virtual Machine	Inter VM and public network
iscsi-vlan10	eth2 on ora2, ora3	Storage Virtual Machine	For iscsi vlan 10
iscsi-vlan20	eth3 on ora2, ora3	Storage Virtual Machine	For iscsi vlan 20

¹The management network was automatically created during the installation. Create the other networks after the installation.

3.4 Create virtual machines for Oracle RAC

Virtual machines can be created fresh or imported from predefined VM templates. VM templates can be a great way to deploy a new environment if the software included in the template meets the business requirement. Oracle provides many VM templates that can be downloaded. For more information, see the page, [Oracle VM Templates](#).

To create new VMs, navigate to the Servers and VMs tab in Oracle VM Manager and click the Create Virtual Machine icon on the left pane. The following settings were used for the VMs. They should be reviewed and adjusted according to the business requirement.

Parameters	RAC node 1 VM	RAC node 2 VM
Name	vora10	vora11
Server Pool	TechSol	TechSol
Server ¹	ora2	ora3
Repository	TechSolRepo2	TechSolRepo2
Enable High Availability	Unchecked	Unchecked
Operating System	Oracle Linux 7	Oracle Linux 7
Mouse Device Type	OS Default	OS Default
Keymap	en-us	en-us
Domain Type ²	XEN_HVM_PV_DRIVERS	XEN_HVM_PV_DRIVERS
Start Policy	Best Server	Best Server
Max Memory (MB)	8000	8000
Memory (MB)	8000	8000
Max. Processors	1	1
Processors	1	1
Priority	50	50
Processor Cap%	100	100
Network ³	VMNetwork-VM VMNetwork-Oracle-private iscsi-vlan10 iscsi-vlan20	VMNetwork-VM VMNetwork-Oracle-private iscsi-vlan10 iscsi-vlan20
Disks ⁴	CD/DVD Virtual Disk (50GB) Use Sparse Allocation	CD/DVD Virtual Disk (50G) Use Sparse Allocation
Boot Order	CDROM Disk	CDROM Disk

¹ Server: Each VM was attached to a different server. The RAC nodes should run on different server if possible.

² Domain Type: XEN_HVM_PV_DRIVERS type is full hardware virtualization type with additional paravirtualized drivers for improved performance. Other domain types are XEN_HVM or XEN_PVM. Which domain type to use depends on the guest OS. For Oracle Linux 7, only XEN_HVM and XEN_HVM_PV_DRIVERS are supported.

³ Network: Create one Virtual NIC for each network and use dynamically assign MAC for the Virtual NIC.

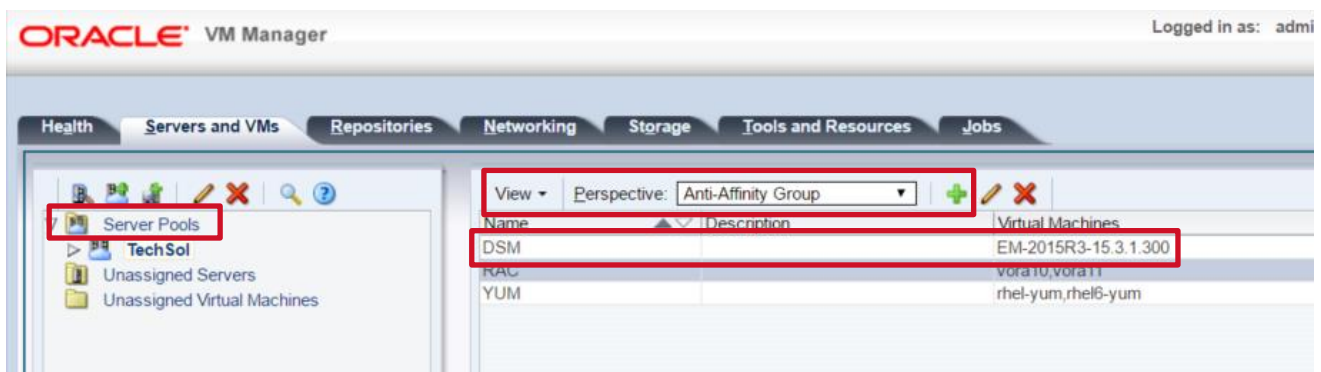
⁴ Disks: Assign a CD/DVD slot and point to the OS ISO file in the repository. Also create a virtual disk to hold the OS

For more information on these settings, see Oracle VM documentation at [Oracle Virtualization Product Documentation Libraries](#).

3.5 Create anti-affinity group

To achieve load balancing and redundancy for the RAC environment, the two Oracle RAC VMs are put into an anti-affinity group which prevents them from running on the same Oracle VM Server.

To create an anti-affinity group in Oracle VM Manager, navigate to the **Servers and VMs** tab and select the server pool. In the **Perspective** drop-down menu, select **Anti-Affinity Group** and click the **+** icon. Enter a name and select **vora10** and **vora11** to the group.



3.6 Install the guest OS

The guest OS installation is performed in the VM console. The installation process is no different than a normal Oracle Linux installation. The selection of the base environment software depends on the user and application requirement. If migrating existing RAC database to the VMs, the software should match the original servers as closely as possible. The following example shows a software selection that works well for an Oracle RAC configuration.

Base environment

- Oracle Linux 7 Server with GUI

Add-ons

- Network File System Client
- File and Storage Server

Additional utility packages

- ntpdate
- ntp

- sg3_utils-libs
- sg3_utils
- iscsi-initiator-utils
- dm-mapper-multipath

Packages required by Oracle 12c:

- xorg-x11-xauth
- xorg-x11-utils
- libstdc++-devel
- libaio-devel
- ksh
- gcc-c++
- compat-libstdc++-33
- compat-libcap1

All Oracle UEK kernels now provide Paravirtualized Drivers natively and there is no need to install the extra add-ons separately.

3.7 **Configure VM network adapter**

In the guest OS, configure the four virtual interfaces specified in section 3.4. Define IPs, netmask, gateway, DNS, and time server using the normal procedures for Oracle Linux. The assigned IP addresses should correspond to the same IP subnets of the Oracle VM networks.

To identify the virtual interface for a specific network, compare the MAC address shown in the guest OS to the list of **VNICs** in the **Network** information pane of the VM.

The screenshot displays the Oracle VM configuration interface. The top section shows a list of virtual machines with columns for Name, Status, Tag(s), Event Severity, Server, Max. Memory (MB), Memory (MB), Max. Processors, and Processor. Below this, the 'Configuration' tab is active, showing a table for network settings. A blue callout box labeled 'VNIC mapping' points to the 'VMNetwork-Oracle-private' entry in the table. Below the configuration table, a terminal window shows the output of the 'ip addr show' command, highlighting the network interface 'eth0' with its IP address '172.16.25.247' and MAC address '00:21:f6:4a:78:43'.

Name	Status	Tag(s)	Event Severity	Server	Max. Memory (MB)	Memory (MB)	Max. Processors	Processor
EM-2015R3-15.3...	Running		Informational	ora2	16000	16000	2	2
EM-APM.0	Running		Informational	ora2	8000	8000	2	2
OEM12c	Stopped		Informational	ora2	16384	16384	2	2
Test	Stopped		Informational	ora3	1024	1024	1	1
rhel6-yum	Running		Informational	ora2	2048	2048	1	1
rhel-yum	Running		Informational	ora3	2048	2048	1	1
vora10	Running		Informational	ora2	8000	8000	1	1

VNIC	Ethernet Network	IP Addresses
00:21:f6:4a:78:43	VMNetwork-Oracle-private	172.16.25.247, fe80::221:f6ff:fe4a:7843
00:21:f6:86:c5:9f	VMNetwork-Oracle-private	10.10.90.11, fe80::221:f6ff:fe86:c59f
00:21:f6:50:12:98	iscsi-vlan10	10.10.90.11, fe80::221:f6ff:fe50:1298
00:21:f6:7c:44:95	iscsi-vlan20	10.20.90.11, fe80::221:f6ff:fe7c:4495

```

[root@vora10 grub2]# ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
    valid lft forever preferred lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qdisc pfifo_fast
    link/ether 00:21:f6:4a:78:43 brd ff:ff:ff:ff:ff:ff
    inet 172.16.25.247/20 brd 172.16.31.255 scope global eth0
    inet 172.16.25.248/20 brd 172.16.31.255 scope global secondary eth0:1
    inet6 fe80::221:f6ff:fe4a:7843/64 scope link
    valid lft forever preferred lft forever
  
```

For the 10G network, follow the Oracle knowledge article 1519875.1 to adjust the TCP settings that improves the network performance. The article can be obtained on [My Oracle Support](#) website settings.

The recommended settings for 10GB network are listed in Table 16. To make them persistent across system reboots, add them to the `/etc/sysctl.conf` file.

Table 16 Oracle recommended TCP settings for 10GB network

Parameter	Value	Description
net.core.rmem_max	134217728	Maximum of the receive buffer size used by each TCP socket
net.core.wmem_max	134217728	Maximum of the send buffer size used by each TCP socket
net.ipv4.tcp_rmem	4096 87380 134217728	Auto-tune TCP buffer limits. Min, Default, Max size of the receive buffer used by each TCP socket
net.ipv4.tcp_wmem	4096 65536 134217728	Auto-tune TCP buffer limits. Min, Default, Max size of the send buffer used by each TCP socket
net.core.netdev_max_backlog	300000	Maximum number of incoming connections backlog queue
net.ipv4.tcp_moderate_rcvbuf	1	Auto-tune the receiver buffer size

net.bridge.bridge-nf-call-iptables	0	netfilter
net.bridge.bridge-nf-call-arptables	0	netfilter
net.bridge.bridge-nf-call-ip6tables	0	netfilter

3.8 Configure iSCSI initiator

Configure the iSCSI initiator to access the SC Series array. Here is a summary of the steps:

1. Set the following values in the `/etc/iscsi/iscsid.conf` file:

```
node.session.timeo.replacement_timeout = 5
node.session.cmds_max = 1024
node.session.queue_depth = 128
```

2. Perform iSCSI discovery against the SC Series array. 10.10.7.10 and 10.20.7.10 are the iSCSI fault domain control port IPs.

```
# iscsiadm -m discovery -t st -p 10.10.7.10:3260
# iscsiadm -m discovery -t st -p 10.20.7.10:3260
```

3. Log in to the discovered iSCSI qualified name (iqns).

```
# iscsiadm -m node -login
```

For more information about configuring the iSCSI initiator on Oracle Linux, refer to the documents, [Dell SC Series Storage Best Practices for Oracle VM](#) and [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 7x Best Practices](#).

3.9 Configure multipathing and udev rules

With dual iSCSI connections available on the VMs, set up multipathing to provide path failover and load balancing. For database volumes, proper permission must also be set on the devices. The following example shows using udev and multipath to achieve persistent settings on a device.

The following line is added in the `/etc/udev/rules.d/99-oracle-asmdevices.rules` file. The rule applies to any multipath devices with a prefix of ORA.

```
KERNEL=="dm-*", ENV{DM_NAME}=="ORA*", OWNER="grid", GROUP="dba", MODE="0660"
```

Create a multipath alias for each database volume in `/etc/multipath.conf` using the `wwid` identifier. The following example sets the multipath device name to `ORA_GIDATA`. The udev rule above then sets the ownership and permission to `grid:dba` and `660`.

```
multipath {
    uid 200
    gid 200
    wwid "36000d31000fcaf0000000000000000020"
    alias ORA_GIDATA
```

```
}
```

Reload the multipath configuration after changes are made in the `/etc/multipath.conf` file:

```
# systemctl reload multipathd
```

3.10 Add VMs to DSM

Once the iSCSI initiator is set up in section 3.8, the VM iqns, along with their IP addresses, become available to DSM. Add each VM to DSM using the iqns.

4 Install Oracle Grid Infrastructure and DB software

Install the Oracle Grid Infrastructure and RDBMS software in both VMs. The installation location can be the virtual disks that hold the guest OS since each VM holds a local copy of the software.

Grid Infrastructure requires a shared crs/voting disk for all the nodes. An iSCSI SC volume is mapped to both VMs for this purpose. Table 17 shows the volume attributes.

Table 17 SC volumes for Grid Infrastructure and RAC database

Volume attributes	CRS/Voting	Database Volumes ¹	Archive Volumes ¹
Name	crsvol	datavol1, datavol2, datavol3, datavol4	fravol1, fravol2
Size	10G	10G each	10G each
Snapshot Profiles	Daily	RACCG ²	RACCG ²
Storage Profile	Recommended (All Tiers)	Recommended (All Tiers)	Recommended (All Tiers)
Read/Write Cache	Enabled/Disabled	Enabled/Disabled	Enabled/Disabled
Compression	Disabled	Disabled	Disabled

¹The database and archive volumes are required only if a new database is being created.

²A Consistent Snapshot Profile is applied to the database and archive volumes to ensure snapshots taken on the database is consistent and recoverable.

4.1 Perform Oracle DB preinstallation tasks

The easiest way to preconfigure the guest OS for Oracle Database 12c or 11g is to download the **oracle-rdbms-server-12cR1-preinstall** or **oracle-rdbms-server-11gR2-preinstall rpm** package. The rpm package performs preinstallation tasks including:

- Update kernel parameters in `/etc/sysctl.conf`
- Download and install required software dependencies
- Create the user, oracle and groups, oinstall and dba
- Configure oracle user limits in `/etc/security/limits.conf`

The rpm package can be downloaded from the [Oracle public yum repository](#). By default, Oracle Linux installs a public yum repository configuration file in `/etc/yum/repos.d`. Make sure the repository entries are enabled in the file.

To install the rpm package and run the preconfiguration tasks:

```
# yum install oracle-rdbms-server-12cR1-preinstall
```

For more information on the rpm packages, see the Oracle article, [How I Simplified Oracle Database 12c and 11g Installations on Oracle Linux 6](#).

Note that the rpm package only creates the user, oracle. If a separate user is required to manage the Oracle Grid Infrastructure, create the user grid manually and duplicate the user settings from the user, oracle.

Set up ssh keys for both oracle and grid user and configure passwordless ssh connectivity across all cluster nodes. The Oracle installer checks for the ssh connectivity as part of the prerequisite.

4.2 Install the Grid Infrastructure and RDBMS software

It is recommended to follow the official installation procedures of Oracle Grid Infrastructure and RDBMS software. The official installation guides can be found in the [Oracle Database Documentation Library](#).

The parameters listed in Table 18 are used for the Grid Infrastructure installation.

Table 18 Grid Infrastructure installation parameters

Parameters	Values	Comments
ORACLE_BASE	/u01/app/grid	
ORACLE_HOME	/u01/app/12.1.0/grid	
OSDBA_GROUP	oinstall	
OSASM_GROUP	dba	
Cluster scan name	voracust1-scan	This must be resolvable in DNS or in /etc/host file
Cluster type	Standard	
Cluster name	voracust1	
Configure GNS	false	
Cluster nodes	vora10-vip.techsol.local vora11-vip.techsol.local	The node vip names must be resolvable in DNS or in /etc/host file
Network Interface Type	eth0 – Public eth1 – Private eth2, eth3 – Do not use	
Storage Type	ASM	
IPMI	False	
ASM disk group	GIDATA	
ASM disk group redundancy	EXTERNAL	When using external SAN storage, EXTERNAL redundancy should be used
ASM AUSize	1MB	
ASM disk group disks	/dev/mapper/ORA_GIDATA	This is the crs/voting disk.
ASM disk discovery string	/dev/mapper/	

The settings listed in Table 19 are used for the Oracle RDBMS software installation.

Table 19 Oracle RDBMS software installation parameters

Parameters	Values
Inventory location	/u01/app/oraInventory
Languages	en
ORACLE_BASE	/u01/app/oracle
ORACLE_HOME	/u01/app/oracle/product/12.1.0/dbhome_1
Database Edition	Enterprise Edition
DBA_GROUP	dba
OPER_GROUP	dba
BACKUPDBA_GROUP	dba
DGDBA_GROUP	dba
KMDBA_GROUP	dba
Cluster nodes	vora10, vora11

4.3 Create ASM disk groups

ASM is the recommended storage type for Oracle RAC databases. Create a minimum of one ASM disk group to hold the data files and archive logs. After the Grid Infrastructure software is installed, the asmca GUI utility is available to create and manage ASM disk groups. The ASM disk groups shown in Table 20 are established for a new RAC database.

Table 20 ASM disk groups for RAC database

ASM disk group name	Type	Total size	Sector size	Block size	AU	Disk Name
DEMODB_DATADG	EXTERNAL	40G	512	4096	1M	/dev/mapper/ORA_DEMODB_DATA1 /dev/mapper/ORA_DEMODB_DATA2 /dev/mapper/ORA_DEMODB_DATA3 /dev/mapper/ORA_DEMODB_DATA4
DEMODB_FRADG	EXTERNAL	20G	512	4096	1M	/dev/mapper/ORA_DEMODB_FRA1 /dev/mapper/ORA_DEMODB_FRA2

4.4 Create a RAC database

Launch the dbca GUI utility from the RDBMS software directory to create a new RAC database. The settings listed in Table 21 are used.

Table 21 RAC database installation parameters

Parameters	Values
Inventory location	/u01/app/oralInventory
Database type	RAC
Cluster nodes	vora10, vora11
Configuration type	Admin managed
Starter Database Type	General Purpose
Database name	vdemodb
Container DB	False
Character set	AL32UTF8
Automatic Memory Management	True
Memory Limit	2G
Enable recovery	True
Storage type	ASM
ASM disk group	DEMODB_DATADG, DEMODB_FRADG

5 Migrate Oracle RAC database from physical servers to VMs

Traditionally, migrating an Oracle RAC database between servers involves copying/moving data from one server to another or using some type of backup/restore solutions such as Oracle RMAN. These approaches often take a large amount of time as the data is being moved and copied across the network or to and from backup mediums. It also poses a potential performance impact to the original database while the copy is taken place.

The SC Series array offers a storage-based alternative to these traditional methods. The block-level snapshot feature, also known as Replay, provides a quick and efficient way to copy or migrate an Oracle RAC database from physical servers to VMs. Taking a snapshot is fast and incurs minimum overhead on the servers. Multiple snapshots can be taken over time. Each snapshot represents a point-in-time copy of a volume and can be used as a recovery point in the event of data corruption or loss due to events such as hardware failure or human errors.

When snapshots are created, the data blocks are frozen (they can be read but not modified). Changes to the existing data are written to new blocks and take precedence over the frozen blocks.

Using snapshots provides the following benefits:

- Creates space-efficient, time-based snapshots at any time for the entire database
- Takes only a few seconds to create a snapshot
- Creates a point-in-time copy of a database on new servers or VMs quickly without moving or duplicating the data
- Allows you to practice migrating the database without impacting the original database
- Provides extra protection to the database; in the event the migration failed, the database can be quickly reverted back to the original volumes

To use snapshots to migrate a database, the following requirements must be met:

- Both original servers and new VMs must have access to the same SC volumes.
- If the original servers and new VMs are on two different SC Series arrays, replication must be set up between them to replicate the database volumes to the second array.
- To take a snapshot of a running database, the number of database volumes must not exceed 40. This is a limitation on a Consistent Snapshot Profile.
- The database must be put in hot backup mode or shut down before taking a snapshot.
- The hardware platform and OS must be compatible between the original servers and new VMs.

5.1 Steps to migrate a RAC database

In this example, a two-node Oracle RAC database is running on two PowerEdge servers. The servers have access to an SC Series array on an FC SAN. The new VMs to which the database is migrated to also have access to the same SC Series array on an iSCSI SAN.

4. Prepare the VMs:
 - a. Follow the steps in section 3 to prepare the new VMs.
 - b. Create a new RAC cluster with compatible software from the original servers. It is not necessary to create a new database. See section 4.
5. Create snapshots on the database volumes in DSM:
 - a. Create a Consistent Snapshot Profile:
 - i. Add all database volumes to a Consistent Snapshot Profile. The profile ensures that IOs are halted across all volumes as a group and snapshots are taken on all volumes simultaneously. This is a requirement for taking a consistent database snapshots.

This step can be skipped if the database is to be shut down before taking a snapshot.

- ii. In DSM, navigate to the **Storage** tab, **Snapshot Profiles** tree. Right-click the profile and select **Create Snapshot Profile**.

The following screen shows a Consistent Snapshot Profile, MARKETDB_CG, consisting of four database volumes.

Name	Expiration	Frequency	Time	Additional Information
OnDemand	1 hour	Once	Date 4/25/16 10:57 AM	

Name	Configured Space	% Full	Active Space	Snapshot Overhead	Actual Space	Total Disk Space	Est. Data Reduction
MARKETDB_DATA 1	500 GB	0.77%	3.85 GB	4.21 GB	8.06 GB	9.14 GB	
MARKETDB_DATA 2	500 GB	0.77%	3.85 GB	4.3 GB	8.15 GB	9.24 GB	
MARKETDB_FRA 1	500 GB	0.85%	4.25 GB	230.5 MB	4.47 GB	5.04 GB	
MARKETDB_FRA 2	500 GB	0.86%	4.29 GB	191 MB	4.47 GB	5.03 GB	

- b. If the database is open, put the database in hot backup mode. If it is not possible because archive log mode is not enabled, shut down the database instead.

```
SQL> alter database begin backup
```

c. Create snapshots:

- Right-click the snapshot profile and select **Create Snapshot**.
- Check **Do Not Expire** or set the **Expire Time** to X number of days to keep the snapshots long enough for the migration.
- End database hot backup mode: Take the database out of hot backup mode or restart the database if it was shut down previously.

SQL> **alter database end backup**

6. Create view volumes from snapshots and map to the VMs:

- Available snapshots are listed in the **Snapshots** information pane of each volume. Select a snapshot and click the **Create Volume from Snapshot** icon. Create a view volume from each database volume.

The screenshot shows the Dell EMC SC Series Storage management console. The left sidebar displays a tree view of storage resources, including 'MARKETDB_DATA 1'. The main pane shows the 'Snapshots' tab for this volume. It includes 'Server Usage' and 'Disk Usage' statistics. Below these, a table lists snapshots with columns: Freeze Time, Expire Time, Size, % of Actual, Description, Space Recovery, Consistent, Source, and Create Volume. The snapshot '4/25/16 10:58:59 AM' is highlighted, and the 'Create Volume from Snapshot' button is circled in red.

Freeze Time	Expire Time	Size	% of Actual	Description	Space Recovery	Consistent	Source	Create Volume
Active		15.5 MB	0.2%		No	No	System	MARKETDB_DATA 1
4/25/16 12:30:44 PM	Never Expires	19.5 MB	0.25%	Created by APM	No	Yes	Consistency Group	MARKETDB_DATA 1
4/25/16 12:15:46 PM	Never Expires	19.5 MB	0.25%	Created by APM	No	Yes	Consistency Group	MARKETDB_DATA 1
4/25/16 12:00:53 PM	Never Expires	96 MB	1.23%	Created by APM	No	Yes	Consistency Group	MARKETDB_DATA 1
4/25/16 11:15:46 AM	Never Expires	94.5 MB	1.21%	Created by APM	No	Yes	Consistency Group	MARKETDB_DATA 1
4/25/16 10:58:59 AM	Never Expires	13 MB	0.17%	Manually Created	No	Yes	Consistency Group	MARKETDB_DATA 1
4/25/16 10:45:46 AM	Never Expires	24 MB	0.31%	Created by APM	No	Yes	Consistency Group	MARKETDB_DATA 1

- Create multipath alias definitions for the view volumes in the `/etc/multipath.conf` file on both VMs. See section 3.9.
- Map the view volumes to both VMs.
- ssh to the VMs as root and scan for new volumes.

```
# rescan-scsi-bus.sh -a
```

7. Mount the ASM disk groups as user grid on both VMs.
The view volumes retain the same ASM disk information as the original volumes because they are exact replicas. ASM on the VMs automatically recognizes these new volumes for disk group MARKETDB_DATA and MARKETDB_FRA.

```
grid$ export ORACLE_SID=+ASM1
grid$ asmcmd mount MARKETDB_DATA
grid$ asmcmd mount MARKETDB_FRA
grid$ asmcmd lsdg -discovery
```

8. Generate a copy of the database parameter file from the original database and transfer it to the new VMs.

```
SQL> create pfile='/tmp/orclmkdb.pfile' from spfile;
oracle $ scp /tmp/orclmkdb.pfile vora10:/tmp/orclmkdb.pfile
```

9. On one of the VMs, start up the database as user, oracle, using the pfile:
- If the VMs have different memory settings than the original servers, modify the pfile to reflect suitable settings. For example, adjust **memory_target**, **sga_target** and other memory settings based on available memory on the VMs.
 - Create any directories referred in the pfile ahead of time on all cluster nodes.
 - Start up the database using the modified pfile and perform database recovery.

```
SQL> startup pfile='/tmp/orclmkdb.pfile' mount;
SQL> alter database recovery;
```

- d. Open the database.

```
SQL> alter database open;
```

10. Create a new spfile from the modified pfile.

```
SQL > create spfile='+MARKETDB_DATA/ORCLMKDB/PARAMETERFILE/newspfile' from
pfile='/tmp/orclmkdb.pfile' ;
```

11. Record the new spfile location from step 10 in the instance pfile on each VM. The instance pfile is located in \$ORACLE_HOME/dbs/\$DB_INSTANCE.pfile.

```
spfile='+MARKETDB_DATA/ORCLMKDB/PARAMETERFILE/newspfile'
```

12. Allow users to access the database to validate the data.

13. Create cluster resources for the database as user, oracle, on the VMs.
- On the original server, run `srvctl` to extract the resource information.

```
oracle$ srvctl config db -d orclmkdb
Database unique name: orclmkdb
Database name: orclmkdb
Oracle home: /u01/app/oracle/product/12.1.0/dbhome_1
Oracle user: oracle
Spfile: +MARKETDB_DATA/ORCLMKDB/PARAMETERFILE/spfile.273.904573853
Password file: +MARKETDB_DATA/ORCLMKDB/PASSWORD/pwdorclmkdb.256.904573571
Domain:
Start options: open
Stop options: immediate
Database role: PRIMARY
Management policy: AUTOMATIC
Server pools:
Disk Groups: MARKETDB_DATA, MARKETDB_FRA
Mount point paths:
Services: orclmkdbsvc
Type: RAC
Start concurrency:
Stop concurrency:
OSDBA group: dba
OSOPER group:
Database instances: orclmkdb1, orclmkdb2
Configured nodes: ora6, ora7
Database is Policy managed
```

- On the VMs, create cluster resources based on the information from step a. Replace the `$FIELD` with values from above, except for the `$SPFILE` which has a new value as defined in step 10.

```
oracle$ srvctl add database -d orclmkdb -o $ORACLE_HOME -spfile '$SPFILE' \
\-pwfile '$PWFILE'

oracle$ srvctl add instance -d orclmkdb -i orclmkdb1 -n vora10
oracle$ srvctl add instance -d orclmkdb -i orclmkdb2 -n vora11
```

This process can be repeated again by DBAs and storage administrators to practice the migration steps ahead of the actual migration. During the actual migration, shut down the original database (step 5) before taking a snapshot. Once the database is migrated and approved by the users, clean up the original volumes by unmapping them from the original servers and then deleting them from DSM.

1 Terminology

Assembly: Contains multiple virtual machine configurations including virtual disks and connectivity settings between them. It simplifies creating a group of related virtual machines.

Data Progression: An SC Series array feature that moves pages between tiers and drive types, as well as among multiple RAID levels within the same tier.

Dynamic capacity: An SC Series array feature that provides block-level storage thin provisioning.

Fault domain: Fault domain must be created for Virtual Ports. FC ports and iSCSI ports cannot mix in a same fault domain. Virtual ports can fail over to other physical ports in the same domain.

Hypervisor: Part of an Oracle VM Server. It is an abstraction layer responsible to manage and monitor all the hardware resources and virtual machine guests. On the Linux x86 platform, an open-source Xen hypervisor is included. On the SPARC platform, the SPARC hypervisor is built into the firmware.

Multipath: Use of multiple physical disk access paths between storage and servers. Multipath is a requirement for Oracle VM to discover SAN disks.

OCFS2: Oracle Clustered File System (OCFS2) is used by Oracle VM to provide clustered and shared storage to Oracle VM Servers and virtual machines.

Oracle VM Agent: Runs on an Oracle VM Server to facilitate communication between the Oracle VM Server and the Oracle VM Manager.

Oracle VM Manager: A web-based management platform that manages and monitors all Oracle VM Servers and virtual machines.

Oracle VM Server: Consists of an Oracle VM agent and a hypervisor which has direct access to the hardware resources.

Snapshot: Formerly referred to as a Replay, a snapshot is an SC Series array feature that provides a storage point-in-time copy of data.

Snapshot profile: Defines the frequency and when snapshots can run.

Storage connect plug-in: Provides storage-specific functionalities for different type of storage.

Storage profile: Defines which tier accepts initial writes and how Data Progression moves pages between tiers in an SC Series array.

Storage repository: Provides storage space to hold different types of files used by virtual machines including assemblies, virtual machine configuration files, virtual machine templates, CD/DVD media files, and virtual disks.

Virtual machine: A guest operating system that runs within an Oracle VM Server.

Virtual machine template: A fully installed and configured virtual machine with all the required applications and software stacks preinstalled. The template is reusable for creating new virtual machines with the exact image and configuration.

Virtual disk: Provides storage space for the OS and applications in a virtual machine. Multiple virtual disks can be assigned to a virtual machine.

Virtual port: In Virtual Port mode, all front-end ports are utilized. A virtual port can fail over to other physical ports within the same fault domain when a port or controller fails.

View volume: A view volume is a volume created from a previous snapshot in an SC Series array.

WebLogic server: Oracle VM Manager runs as an application within the WebLogic server to provide the application layer and web layer function.

Zone: Fibre Channel zones are used to segment the fabric to restrict access.

1 Configuration files

1.1 Volume group definition for boot volume on Oracle VM Manager server

The volume group is defined with the following specifications:

```
--- Volume group ---
VG Name                vg00
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No   5
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 3
Open LV                 3
Max PV                 0
Cur PV                 1
Act PV                 1
VG Size                 99.51 GiB
PE Size                 4.00 MiB
Total PE                25474
Alloc PE / Size         25474 / 99.51 GiB
Free PE / Size          0 / 0
VG UUID                 hSXw8P-jar8-HWG4-SIB5-7dAo-FT2d-CcZyET

--- Logical volume ---
LV Path                 /dev/vg00/lv_root
LV Name                 lv_root
VG Name                 vg00
LV UUID                 iHfXYO-amlV-b4Ut-tHhE-aIYF-IPQF-qNEhbL
LV Write Access         read/write
LV Creation host, time  ora4, 2013-06-28 08:48:44 -0500
LV Status                available
# open                  1
LV Size                 50.00 GiB
Current LE              12800
Segments                1
Allocation              inherit
Read ahead sectors      auto
- currently set to      256
Block device            253:7
```

--- Logical volume ---

LV Path /dev/vg00/lv_home
LV Name lv_home
VG Name vg00
LV UUID SNpmUb-P3dw-voCM-ITKp-UwcF-eTUw-Dw5bEB
LV Write Access read/write
LV Creation host, time ora4, 2013-06-28 08:48:50 -0500
LV Status available
open 1
LV Size 18.01 GiB
Current LE 4610
Segments 1
Allocation inherit
Read ahead sectors auto
- currently set to 256
Block device 253:9

--- Logical volume ---

LV Path /dev/vg00/lv_swap
LV Name lv_swap
VG Name vg00
LV UUID ii7FWp-9H2h-TDIR-CcAw-4DR2-rYoO-2ae7bq
LV Write Access read/write
LV Creation host, time ora4, 2013-06-28 08:48:55 -0500
LV Status available
open 1
LV Size 31.50 GiB
Current LE 8064
Segments 1
Allocation inherit
Read ahead sectors auto
- currently set to 256
Block device 253:8

--- Physical volumes ---

PV Name /dev/mapper/mpathap2
PV UUID Zwyfxn-XsDR-TN6I-pLTj-6c32-a4DO-k1G6vC
PV Status allocatable
Total PE / Free PE 25474 / 0

1.2 Volume group definition for vgoracle

```
--- Volume group ---
VG Name                vgoracle
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No   6
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 1
Open LV                 1
Max PV                  0
Cur PV                 1
Act PV                  1
VG Size                 20.00 GiB
PE Size                 4.00 MiB
Total PE                5119
Alloc PE / Size         5119 / 20.00 GiB
Free PE / Size           0 / 0
VG UUID                 0c04e4-uMWm-Aywb-5mcg-fnFJ-XHMa-IrUiCH

--- Logical volume ---
LV Path                 /dev/vgoracle/lv_u01
LV Name                 lv_u01
VG Name                 vgoracle
LV UUID                 9jcE57-C3Fv-YhxX-1J81-aQQp-dlv3-dkbQBo
LV Write Access         read/write
LV Creation host, time  ora4.techsol.beer.town, 2014-04-05 15:31:23 -0500
LV Status                available
# open                   1
LV Size                  20.00 GiB
Current LE               5119
Segments                 1
Allocation               inherit
Read ahead sectors       auto
- currently set to       256
Block device             253:12

--- Physical volumes ---
PV Name                  /dev/mapper/mpathb
PV UUID                  fCE3Bf-fPu2-UeLp-LMYy-eIav-wsmN-gbP0uR
PV Status                allocatable
Total PE / Free PE       5119 / 0
```

1.3 Multipath configuration on Oracle VM Servers

```
defaults {
    user_friendly_names no
    getuid_callout "/lib/udev/scsi_id --whitelisted --replace-whitespace --
device=/dev/%n"
}
blacklist {
    devnode "^ (ram|raw|loop|fd|md|dm-|sr|scd|st|nbd) [0-9] *"
    devnode "^hd[a-z] [0-9] *"
    devnode "^etherd"
    devnode "^nvme.*"
    %include "/etc/blacklisted.wids"
}
devices {

    device {
        vendor                "IBM"
        product                "1724-100"
        hardware_handler      "1 rdac"
        path_grouping_policy   group_by_prio
        prio                   rdac
        path_checker           rdac
        no_path_retry          10
    }

    device {
        vendor                "IBM"
        product                "1742-900"
        hardware_handler      "1 rdac"
        path_grouping_policy   group_by_prio
        prio                   "rdac"
        failback               immediate
        path_checker           rdac
        no_path_retry          10
    }

    device {
        vendor                "XIV"
        product                "NEXTRA"
        path_grouping_policy   multibus
        rr_min_io              1000
        path_checker           tur
        failback               immediate
        no_path_retry          10
    }
}
```

```

device {
    vendor            "IBM"
    product            "2810XIV"
    path_grouping_policy  multibus
    rr_min_io          1000
    path_checker        tur
    failback            immediate
    no_path_retry       10
}

device {
    vendor            "HP"
    product            "MSA1510i VOLUME"
    path_grouping_policy  group_by_prio
    path_checker        tur
    prio               "alua"
    no_path_retry       10
}

device {
    vendor            "DataCore"
    product            "SAN*"
    path_grouping_policy  failover
    path_checker        tur
    failback            10
    no_path_retry       10
}

device {
    vendor            "EQLOGIC"
    product            "100E-00"
    path_grouping_policy  failover
    failback            immediate
    no_path_retry       10
}

device {
    vendor            "COMPELNT"
    product            "Compellent *"
    path_grouping_policy  multibus
    path_checker        tur
    failback            immediate
    rr_min_io          1024
    no_path_retry       10
}

device {
    vendor            "FALCON"
    product            ".*"
    path_grouping_policy  multibus

```



```

        failback            immediate
            no_path_retry    10
    }
    device {
        vendor                "EMD.*"
        product                "ASTRA (ES 12F) | (SA 16i)"
        path_grouping_policy    failover
        failback                immediate
        path_checker            tur
        no_path_retry            10
    }
device {
    vendor                "FUJITSU"
    product                "E[234]000"
    path_grouping_policy    group_by_prio
    prio                    "alua"
    failback                immediate
    no_path_retry            10
    path_checker            tur
}
device {
    vendor                "FUJITSU"
    product                "E[68]000"
    path_grouping_policy    multibus
    failback                immediate
    no_path_retry            10
    path_checker            tur
}
device {
    vendor                "AC&Ncorp"
    product                "JetStorSAS516iS"
    path_grouping_policy    multibus
    failback                15
    no_path_retry            10
    rr_weight                priorities
    path_checker            tur
}
device {
    vendor                "XYRATEX"
    product                "F5402E | [FE] 5412E | [FE] 5404E | F6512E | [FEI] 6500E"
    path_grouping_policy    failover
    failback                3
    no_path_retry            10
    path_checker            tur
}
device {
    vendor "FUJITSU"

```

```

        product "ETERNUS_DXM|ETERNUS_DXL|ETERNUS_DX400|ETERNUS_DX8000"
        prio alua
        path_grouping_policy group_by_prio
        path_selector "round-robin 0"
        failback immediate
        no_path_retry 10
    }
    device {
        vendor "NETAPP"
        product "LUN.*"
        dev_loss_tmo 50
    }
    device {
        vendor      "ATA*"
        product      ".*"
    }
}

```

1.4 Multipath configuration on Oracle VM Manager server

```

defaults {
    user_friendly_names yes
}
blacklist {
    devnode "(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z]"
    devnode "^dcssblk[0-9]*"
    device {
        vendor "DGC"
        product "LUNZ"
    }
    device {
        vendor "IBM"
        product "S/390.*"
    }
    device {
        vendor "ATA"
    }
    device {
        vendor "3ware"
    }
    device {
        vendor "AMCC"
    }
    device {

```

```

        vendor "HPT"
    }
    device {
        vendor iDRAC
        product Virtual_CD
    }
    device {
        vendor HL-DT-ST
        product DVD-ROM_DU70N
    }

    device {
        vendor "Initio.*"
    }
}
devices {
    device {
        vendor "COMPELNT"
        product "Compellent *"
        path_grouping_policy multibus
        path_checker tur
        failback immediate
        rr_min_io 1024
        no_path_retry 10
    }
}

```

2 Technical support and resources

[Dell.com/support](https://dell.com/support) is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC Storage platforms.

2.1 Related documentation

The [SC Series customer Portal](#) is for existing SC Series customers and requires a valid portal account to access the Knowledge Center. The following documentation can be downloaded from the Knowledge Center:

- Dell Storage Center System Manager Version 6.x Administrator's Guide
- Dell Storage Manager 2016 R1 Administrator's Guide

Dell SC Series papers focused on Oracle include the following:

- [Oracle technical papers and videos](#)
- [Dell SC Series Storage Best Practices for Oracle VM](#)
- [Dell Storage Center with Red hat Enterprise Linux \(RHEL\) 6x Best Practices](#)
- [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 7x Best Practices](#)
- [Dell Networking S5000 Switch Configuration Guide for SC Series iSCSI SANs](#)

The [Oracle Help Center](#) provides useful information and documentation on Oracle products. Oracle provides a vast amount of information on Oracle Virtualization products and support, such as the following documents and resources:

- [Oracle Virtualization product page](#)
- [Oracle VM Server for x86 product page](#)
- [Oracle VM Release 3.3 documentation](#)
- [Oracle virtualization product documentation libraries](#)
- [Oracle Database Documentation library](#)