

Dell EMC SC Series Storage Best Practices for Oracle VM

Henry Wong, Product Specialist
Dell EMC Storage Engineering
July 2017

Revisions

Date	Description
March 2012	Initial draft
March 2013	Initial release
January 2016	Updated with Oracle VM version 3.x, SCOS version 6.x, and use of Dell Enterprise Manager as the preferred storage management tool; refreshed storage connectivity and guidelines; added new section on protecting Oracle VM with Replays
July 2017	Remove references of APM and OEMP

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.

© 2012 - 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] [CML1118]

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
1 Introduction.....	4
1.1 Audience.....	4
1.2 Solution overview.....	4
2 Storage options and connectivity	7
2.1 Fibre Channel SAN.....	7
2.2 iSCSI SAN	12
2.3 Connecting Oracle VM Servers to SC Series.....	16
2.4 Multipathing on Oracle VM Servers.....	17
2.5 Boot from SAN.....	19
3 Oracle VM storage guidelines	21
3.1 Storage connect plug-ins.....	21
3.2 Oracle VM storage repositories and SC Series considerations	21
3.3 Creating Oracle VM storage repository with SC Series	25
3.4 Extending Oracle VM storage repository.....	36
3.5 Oracle VM Manager storage consideration.....	42
4 Protecting Oracle VM with storage snapshots	43
4.1 Data Instant Replay	43
4.2 Taking snapshots of Oracle VM repositories and Oracle VM Manager database	43
4.3 Restoring and recovering Oracle VM files from a snapshot.....	46
4.4 Restoring and recovering Oracle VM Manager	56
1 Terminology.....	59
1 Technical support and resources	61
1.1 Related documentation.....	61

1 Introduction

Dell EMC™ SC Series arrays offer many advanced features such as Dynamic Capacity, Data Progression, and Data Instant Replays. This document provides configuration best practices, examples, tips, recommended settings, and other storage guidelines to better integrate SC arrays with the Oracle® Virtual Machine (VM) virtualization solution.

This document covers the Oracle VM version 3.x for Linux® x86 platform only. It does not cover the installation and administration of Oracle VM software. The official Oracle VM documentation library contains detailed information of typical installation and configuration. For more details, refer to [Oracle Virtualization Product Documentation Libraries](#) and the other resources listed in appendix **Error! Reference source not found..**

Note: This document includes general recommendations that may not be adaptable to all configurations. There are certain circumstances and environments in which the configuration may vary based upon individual or business needs.

1.1 Audience

This is a technical document for system administrators and other information technology professionals who are responsible for the setup and maintenance of Oracle VM for Linux and SC Series arrays.

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 software installations and virtualization concepts

1.2 Solution overview

Oracle VM is a virtualization solution designed to manage and run virtual machines with a variety of applications and workloads. Mainstream guest operating systems are supported such as Oracle Linux, Oracle Solaris, Red Hat Enterprise Linux, and Microsoft® Windows Server®.

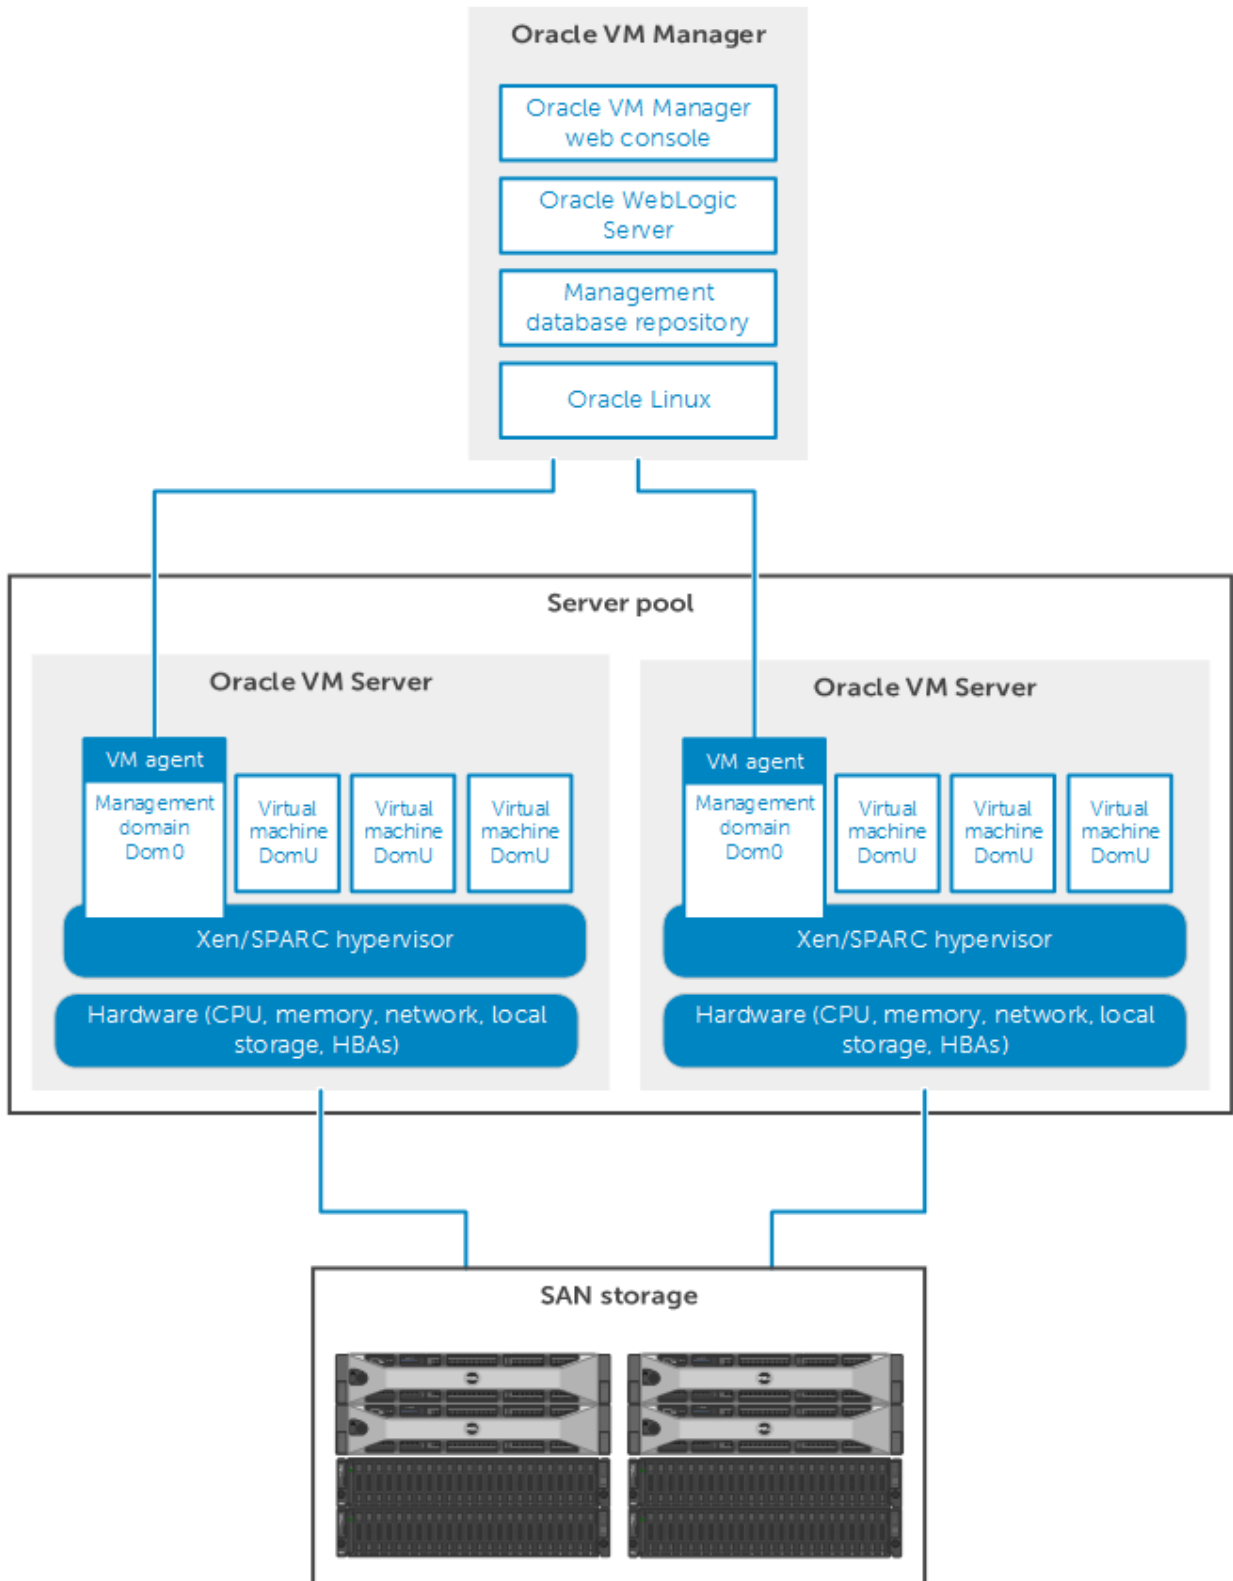


Figure 1 Architecture of Oracle VM

The high-level architecture of the Oracle VM environment (shown in Figure 1) consists of the following components:

Oracle VM Server for x86: This is installed on a 32-bit or 64-bit x86 architecture hardware. The hypervisor is a Xen hypervisor 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 CPU, 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 SPARC: It is similar to the x86 counterpart except that it runs specifically on Oracle SPARC hardware and the SPARC hypervisor is built into the SPARC firmware.

Oracle VM Manager: This is a web-based management tool that runs on Oracle Linux on a separate server. It is built on 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 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 virtual machines running one of the supported operating systems such as Oracle Linux, Red Hat Enterprise Linux, or Microsoft Windows Server. Multiple virtual machines 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.

Shared storage: Third-party NFS, Fibre Channel (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.

For more details, refer to the *Oracle VM Concepts Guide*, available in the [Oracle Virtualization Product Documentation Libraries](#).

2 Storage options and connectivity

Oracle VM supports the following storage types:

- Local disks
- NFS storage
- Fibre Channel SANs
- iSCSI SANs

This section only covers connectivity guidelines for Fibre Channel and iSCSI SANs to Oracle VM Servers and the Oracle VM Manager host.

The guidelines in this document were demonstrated using the Dell SC8000 storage array which supports both Fibre Channel and iSCSI SANs. Depending on the model of the SC Series array, it can support Fibre Channel, iSCSI, or both transport protocols simultaneously. To learn more about the SC Series family, visit the Dell [Storage Solutions](#) page.

2.1 Fibre Channel SAN

For high availability and redundancy, it is recommended to configure the SAN with a minimum of two Fibre Channel switches and two controllers in the SC Series array. Multipathing software is enabled on the servers to manage path failover.




Figure 2 Fibre Channel SAN with dual fabrics, dual controllers, and virtual fault domains

2.1.1 Storage Center virtual ports

Storage Center Operating System (SCOS) introduced the virtual ports feature in version 5.0. With virtual ports, all front-end I/O ports are active. Virtual ports offer many benefits including:

- Increased performance: Virtual ports eliminate the need for reserve ports and therefore free up available ports for either front-end or back-end bandwidth.
- Improved high availability: The virtual port automatically fails over to another physical port within the same domain without bringing down the entire controller.
- Simplified Fibre Channel and iSCSI connectivity: Only the virtual port, WWNs, or IP addresses need to be configured in Fibre Channel zones or iSCSI configuration instead of all physical ports.
- Port-level redundancy within the same controller: This requires a minimum of two connections for each controller.
- Automatic port migration: One virtual port is typically associated to one physical port. When a physical port failed, the virtual port is then migrated automatically to another physical port. Therefore more than one virtual port can be associated to a physical port.
- Virtual fault domains: For each fabric, a separate fault domain must be created. Virtual fault domains determine which physical port the virtual port fails over to in the event of a controller failure. For example, Figure 3 shows two fault domains created on a SC Series array. Each domain consists of four virtual ports and four physical ports span across two controllers. Any virtual ports can fail over to any physical ports in the same domain.

Note: Dell EMC strongly recommends using virtual ports on SC Series arrays.

An SC Series array running in legacy mode can be converted to virtual port mode. This conversion is a one-time operation and it cannot be reversed. Contact Copilot support to assist with this conversion because there might be outage involved (see appendix 1 for Copilot support information). For more information on virtual ports, consult the *Dell Enterprise Manager Administrator's Guide* available on the SC Series [Customer Portal](#) (login required).

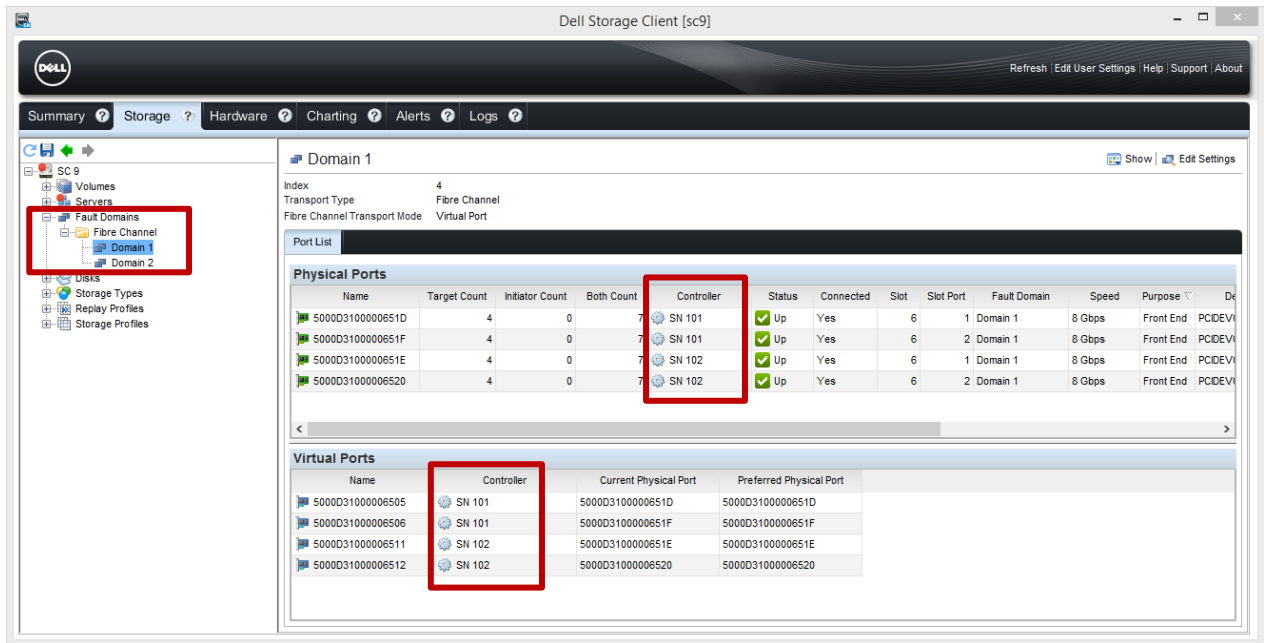


Figure 3 Fault domains and virtual ports in Dell Enterprise Manager

2.1.2 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). Either physical ports (port zoning) on the Fibre Channel switches or the WWNs (name zoning) of the end devices can be used in zoning. It is recommended to use name zoning and SC virtual ports. It allows better flexibility because it is not tied to a specific physical ports on either the switch or the array.

Zoning Fibre Channel switches for Oracle VM hosts is essentially no different than zoning any other hosts to the SC Series arrays. The following list includes key points and options.

Zoning rules and recommendations:

- Name zoning using WWNs is recommended with virtual ports.
- Create a zone that includes all virtual WWNs on the SC Series array controllers. If dual controllers are used, all virtual WWNs on both controllers should be in the same zone.
- Create a second zone that includes only the physical WWNs on the SC Series array controllers.
- For server zones, it is a best practice to use single-initiator, multiple-target zones. For example, for each Fibre Channel port on the server, create a server zone that includes the HBA WWN and all the virtual WWNs on the SC Series array controllers. If the server has four HBA ports, the following zones would be created:
 - One SC virtual port zone
 - One SC physical port zone
 - Four server HBA zones

Note: It is recommended to use name zoning, SC virtual ports, and single-initiator, multiple-target zones.

2.1.3 Configuring Fibre Channel host bus adapter

The following settings should be reviewed and adjusted accordingly on the Oracle VM Servers and the Oracle VM Manager host.

2.1.3.1 QLogic Fibre Channel card BIOS settings

The following BIOS settings are set by the Fast!UTIL BIOS utility, which can be accessed during the system boot process. Make sure the settings are adjusted on all adapter ports.

```
QLE2562 PCI3.0 Fibre Channel ROM BIOS Version 3.20
Copyright (C) QLogic Corporation 1993-2013. All rights reserved.
www.qlogic.com

Press <CTRL-Q> or <ALT-Q> for Fast!UTIL
Firmware Version 5.09.00
```

Table 1 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

2.1.3.2 Emulex Fibre Channel card BIOS settings

The Emulex BIOS settings are set by the LightPulse utility, which can be accessed during the system boot process. Make sure the settings are adjusted on all adapter ports.

```
Emulex LightPulse FC x86 BIOS, Version 10.6.109.0
Copyright (c) 1997-2015 Emulex. All rights reserved.

Press <Alt E> or <Ctrl E> to enter Emulex BIOS configuration
utility. Press <s> to skip Emulex BIOS
Emulex BIOS is Disabled on Adapter 01
Emulex BIOS is Disabled on Adapter 02
Emulex FC BIOS is not installed!!!
```

Table 2 Emulex BIOS settings

Parameter	Value
Topology (point to point only)	2

2.1.3.3 Configure host bus adapter queue depth and module settings

When configuring the host bus adapter on the system, the queue depth might need to be adjusted. The default value 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.

The OS driver module loaded for each HBA ultimately regulates the HBA settings including the queue depth. For example, if the HBA BIOS is set to 255 and the driver module is set to 128, the maximum queue depth for that card or port will be 128.

To set queue depth and other settings:

- QLogic: The module file is located in `/etc/modprobe.d/ql2xxx.conf`. Add the following line to the file:

```
options qla2xxx ql2xmaxqdepth=value
```

- Emulex: The module file is located in `/etc/modprobe.d/lpfc.conf`. Add the following line to the file:

```
options lpfc lpfc_lun_queue_depth=value lpfc_hba_queue_depth=value  
lpfc_devloss_tmo=value
```

`lpfc_devloss_tmo` - Device time out in seconds. Set to 60 seconds in a single path environment or 5 seconds in a multipath environment.

Update the RAMFS disk image:

The `initramfs` file needs to be updated as well. To generate an updated version, use the following command:

```
# mkinitrd -f -v /boot/initramfs-$(uname -r).img $(uname -r)
```

Reboot the system for the change to take effect. After the system rebooted, use `lsscsi -l` command to verify the setting.

```
[root@ora2 ~]# lsscsi -l  
[0:0:0:0] disk ATA MZ-5EA1000-0D3 AXM1 /dev/sda  
state=running queue_depth=31 scsi_level=6 type=0 device_blocked=0 timeout=30  
[4:0:0:0] cd/dvd HL-DT-ST DVD-ROM DU70N D300 /dev/sr0  
state=running queue_depth=1 scsi_level=6 type=5 device_blocked=0 timeout=30  
[6:0:0:1] disk COMPELNT Compellent Vol 0604 /dev/sdb  
state=running queue_depth=128 scsi_level=6 type=0 device_blocked=0 timeout=30  
[6:0:0:2] disk COMPELNT Compellent Vol 0604 /dev/sdc  
state=running queue_depth=128 scsi_level=6 type=0 device_blocked=0 timeout=30  
[6:0:0:6] disk COMPELNT Compellent Vol 0605 /dev/sdd
```

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

2.2 iSCSI SAN

The rules for iSCSI are similar to Fibre Channel. It is recommended to design a storage network with high availability and redundancy. Figure 4 depicts a dual-network, dual-controller, redundant-path setup.

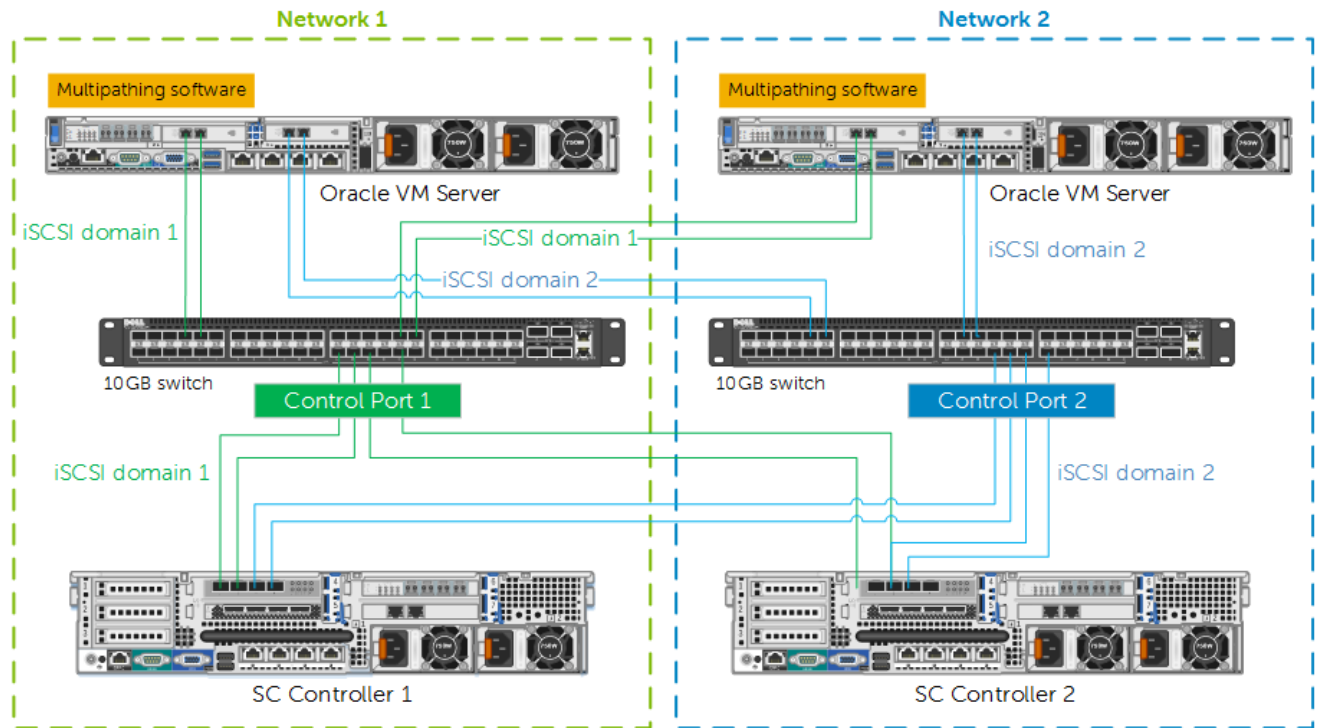


Figure 4 iSCSI SAN with dual networks, dual controllers, control ports, and virtual fault domains

2.2.1 SC Series iSCSI best practices and guidelines

- Use two dedicated iSCSI networks to isolate iSCSI traffic between the Oracle VM Servers and the SC Series array controllers from other public or application traffic.
- Use dual SC Series array controllers and iSCSI fault domains for controller-level redundancy.
- Use dual network switches to provide network-level redundancy.
- Use a minimum of two connections per fault domain from the same controller for port-level redundancy.
- SC virtual ports are strongly recommended.
- A control port is created for each fault domain and only the control port IPs are used in the iSCSI configuration on the Oracle VM Servers. The control port IPs are defined in Dell Enterprise Manager (see Figure 5).
- Servers connect to the control ports which redirect iSCSI traffic to the appropriate physical ports in the same domain.

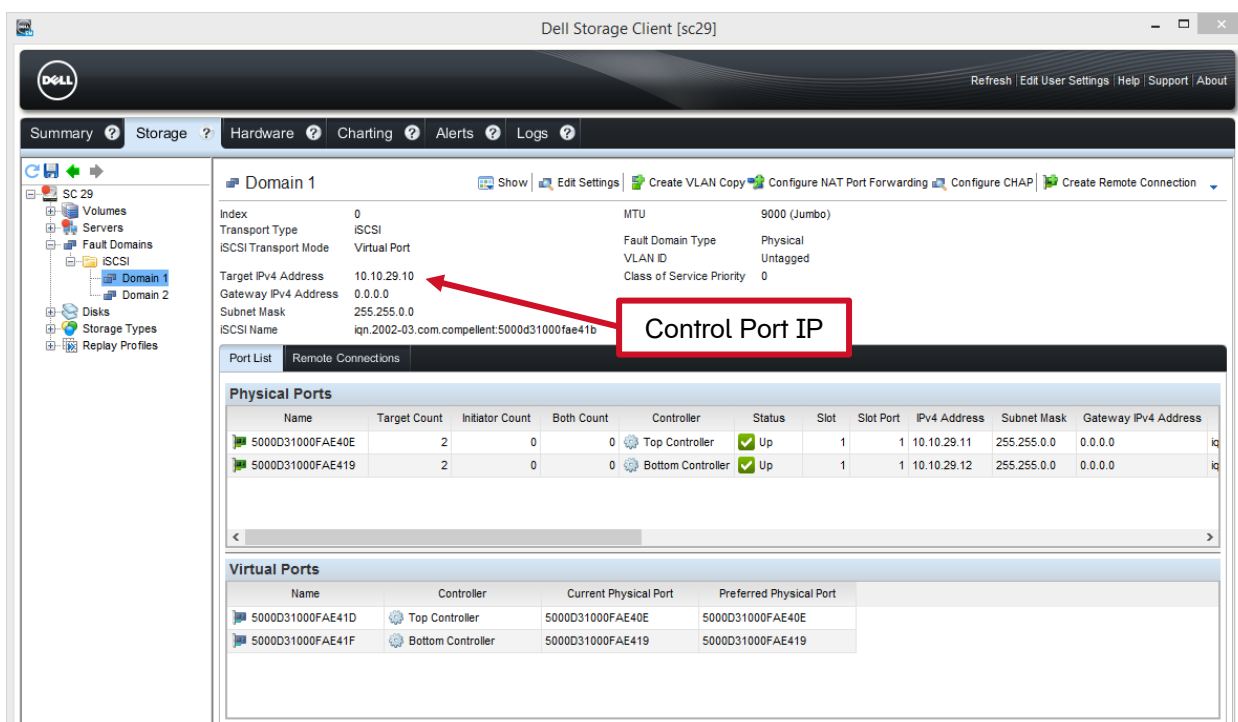


Figure 5 iSCSI fault domain and control port in Dell Enterprise Manager

2.2.2 iSCSI network best practices

- Use a separate virtual local area network (VLAN) or switch for iSCSI traffic.
- For multipathed iSCSI, use two separate VLANs and separate IP subnets.
- Ensure the network bandwidth is sufficient to accommodate all virtual machine needs, with a minimum of gigabit full duplex between SC Series arrays and Oracle VM Servers.
- Disable routing between the regular network and iSCSI network.
- Enable bi-directional flow control on all server and switch ports that handle iSCSI.
- Disable spanning tree on ports connected to the end nodes.

- Consider using Jumbo Frames for improved performance. Jumbo Frames have been supported in Oracle VM Servers since v3.1.1.
- If using Jumbo Frames, make sure all devices on the transmission path have Jumbo Frames enabled and support 9000 maximum transmission unit (MTU).
- Ensure the MTU matches across all end-points including the SC Series array through Dell Enterprise Manager.

2.2.3 Configuring host network adapter and iSCSI

To configure the iSCSI software that is included in Oracle VM Servers, use the following steps to make the recommended settings and activate the iSCSI transport.

1. Edit the `/etc/iscsi/iscsid.conf` file and set the following values. These values dictate the failover timeout and queue depth. The values shown here serve as a starting point and might need to be adjusted depending on the environment.

```
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. This sets up the nodedb in `/var/lib/iscsi` directory. The discovery should be directed to the control port IPs. In the following example, 10.10.29.10 is the control port IP for fault domain 1.

```
# iscsiadm -m discovery -t st -p 10.10.29.10:3260

10.10.29.10:3260,0 iqn.2002-03.com.compellent:5000d31000fae41d
10.10.29.10:3260,0 iqn.2002-03.com.compellent:5000d31000fae41f
```

3. Repeat the command for the second control port.
4. Log in to the discovered iSCSI qualified name (IQN).

```
# iscsiadm -m node -login

Logging in to [iface: default, target: iqn.2002-03.com.compellent:5000d31000fae41f, portal: 10.10.29.10,3260] (multiple)

Logging in to [iface: default, target: iqn.2002-03.com.compellent:5000d31000fae41d, portal: 10.10.29.10,3260] (multiple)

Login to [iface: default, target: iqn.2002-03.com.compellent:5000d31000fae41f, portal: 10.10.29.10,3260] successful.
Login to [iface: default, target: iqn.2002-03.com.compellent:5000d31000fae41d, portal: 10.10.29.10,3260] successful.
```

For more information on configuring iSCSI on the Oracle VM Servers, reference the [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 6x Best Practices](#). The Oracle VM Server is based on Oracle Enterprise Linux 6 which is very similar to Red Hat Enterprise Linux 6.

2.2.4 TCP settings and other tuning considerations

Tuning the TCP buffer sizes can improve the performance significantly especially when 10G network is deployed. The buffer size settings are typically set low by default and they should be carefully reviewed and adjusted accordingly.

For 10 GB network performance tuning, it is recommended to review Oracle knowledge article 1519875.1 available on the [My Oracle Support](#) website. Table 3 summarizes the values recommended in the article. To set these parameters permanently, enter them in the `/etc/sysctl.conf` file and reboot the servers.

Table 3 Oracle recommended TCP settings for 10 GB network


Parameter	Value	Description
net.core.rmem_max	134217728	Maximum receive buffer size used by each TCP socket
net.core.wmem_max	134217728	Maximum send buffer size used by each TCP socket
net.ipv4.tcp_rmem	4096 87380 134217728	Auto-tune TCP buffer limits: min, default, and 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, and 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

2.3 Connecting Oracle VM Servers to SC Series

Once the physical connectivity and proper zoning are in place, the storage connections can be established in Dell Enterprise Manager. A server node definition that defines the server name, operating system type, number of connections, and transport type must be created for each Oracle VM Server and Oracle VM Manager server.

Create a server node definition in Dell Enterprise Manager:

1. In Dell Enterprise Manager, navigate to the **Storage** tab. In the left pane, right-click the **Servers** node tree and select **Create Server**.



5. For iSCSI, the Oracle VM Server IQNs, along with the IP addresses, are listed under the **Host Bus Adapters** section. They become available only after the Oracle VM Server HBA logged into the targets. See section 2.2.3.
For Fibre Channel, the WWNs become available after they are logged into the fabric.
6. Set the server name and operating system type, and select the HBAs. For path redundancy between the servers and storage array, select multiple HBAs.
7. Repeat the process for all Oracle VM Servers in the same cluster. Once the server node definitions are created, the servers and storage are connected.

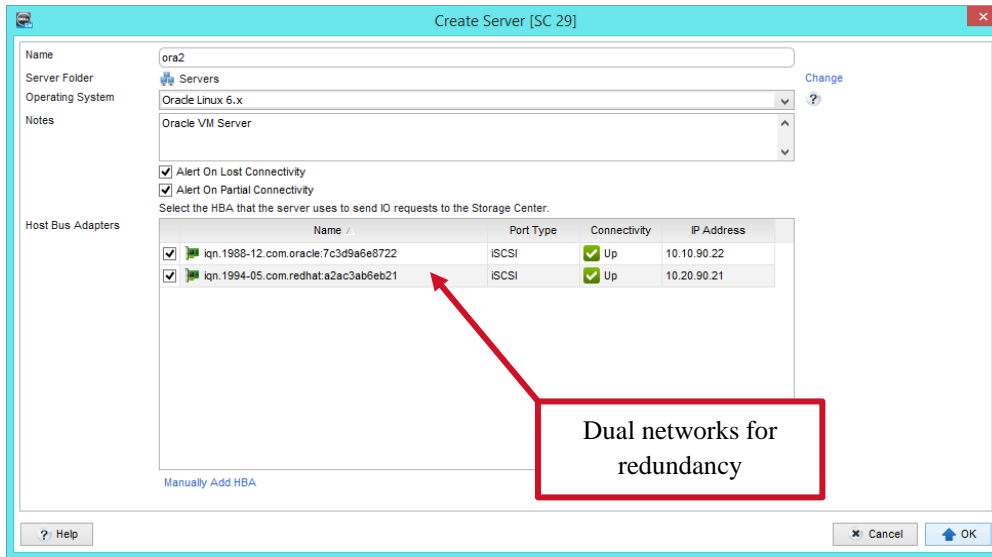


Figure 6 Create server node definition

2.4 Multipathing on Oracle VM Servers

While SC Series arrays provide hardware redundancy and failover with multiple controllers and RAID, multipathing software is used on servers to support and properly utilize multiple connections to the same target. It is an important function that provides load-sharing, load-balancing, and path failover management on the servers. Multipathing is managed at the device level and the multipath devices are accessed through the `/dev/mapper/<DEVICE>` files.

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 for it to discover the SAN disks.

2.4.1 Benefits of using multipath

- DM-Multipath I/O features automatic configuration of the subsystem for a large variety of setups. Active/passive or active/active (with round-robin load balancing) configurations of up to eight paths to each device are supported.
- It takes care of automatic path discovery and grouping as well as automated path retesting, so that a previously failed path is automatically reinstated when it becomes healthy again. This minimizes the need for administrator attention in a production environment.
- Multiple physical connections exist between host bus adapters in the server and the target. Typical connection problems involve faulty adapters, cables, or controllers.

To configure multipathing on the Oracle VM Servers, review and modify the configuration file `/etc/multipath.conf`. The default configuration should work well without changing. Oracle recommends minimal changes to the configuration. The SC Series array specific configuration as follows should already be included in the configuration file under the device section.

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

If changes were made to the configuration, restart the multipathed daemon for it to take effect. The same changes should be made to all Oracle VM Servers.

```
# service multipathd restart
```

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

2.4.2 Showing the status of multipath devices

Multipath device status and information can be shown with the `multipath` command. The following example shows that devices `sdh`, `sdl`, `sdv`, and `sdz` share the same WWN of `36000d3100000650000000000000001f4c`. Therefore, they are the different paths to the same volume on a SC Series array. The volume is 1 TB in size and all paths are in healthy state.

```
# multipath -ll
36000d3100000650000000000000001f4c dm-6 COMPELNT,Compellent Vol
size=1.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
   |- 6:0:1:7 sdh  8:112  active ready running
   |- 6:0:2:7 sdl  8:176  active ready running
   |- 7:0:1:7 sdv  65:80   active ready running
   `-- 7:0:2:7 sdz  65:144 active ready running
```

2.5 Boot from SAN

One object of discussion concerns installing the Oracle VM Server on a multipathed boot drive on a SAN. In some cases, such as with blade servers that do not have internal disk drives, booting from SAN may be the only option, but many Oracle VM Servers can have internal mirrored drives, giving the flexibility to choose. The benefits of booting from SAN are obvious. It alleviates the need for internal drives and allows the ability to replicate or to take snapshots of the boot volume and to migrate the boot volume to another physical system.

Oracle VM 3.1.1 and above supports booting from SAN and it works quite well. However, the decision to boot from SAN depends on many business-related factors including cost, recoverability, and configuration needs. Therefore, Dell does not make specific recommendation on this.

There are a number of things to consider when performing an installation and booting from a SAN volume:

- The Fibre Channel HBA must be configured to boot from a SAN volume. Refer to the adapter manufacturer's documentation on how to perform this task.
- The server BIOS must be configured to select the Fibre Channel HBA first in the boot order.
- The boot volume should have multiple access paths.
- There should be only one boot volume mapped to a server.
- Return to the Fibre Channel HBA and scan for available boot volume. When the boot volume is visible, installation can begin.
- During the Oracle VM Server installation, it will present all local drives and multipathed SAN drives. Make sure to select the correct multipathed drive for the installation.
- It is recommended to use the whole disk for boot volume instead of partitioning the volume.

To create and map a boot volume:

1. Using Dell Enterprise Manager, create a volume and map it to the Oracle VM Servers with the appropriate options enabled.
8. Right-click the volume in Dell Enterprise Manager and click **Map Volume to Server**, select the server, and click the **Advanced Options** link.
9. In the **Advanced Options** screen, select **Map volume using LUN 0** and make sure **Maximum number of paths per Server** under **Configure Multipathing** is set to **Unlimited**. See Figure 7.
10. Click **Finish**.




Figure 7 Map volume advanced option

3 Oracle VM storage guidelines

This section covers configuration, design considerations, and best practices with SC Series array.

3.1 Storage connect plug-ins

SC Series arrays work with Oracle VM using the generic storage connect plug-ins provided by Oracle. The generic plug-ins, both Fibre Channel and iSCSI, are installed by default during the Oracle VM Servers installation and no additional configuration is required for the plug-ins.

With the generic plug-ins, Oracle VM can still benefit from the advanced features provided by SC Series arrays such as volume thin provisioning, Dynamic Capacity, auto storage-tiering, and snapshot technologies.

3.2 Oracle VM storage repositories and SC Series considerations

Storage repositories are used to hold various types of data and files. Oracle VM arranges them into specific folders in each repository. The repositories are accessed through Oracle VM Manager web console in the Repositories tab. Each repository is structured into the following folders:

Assemblies: An assembly contains multiple virtual machine configurations including the virtual disks and connectivity settings between them. It is a reusable template that simplifies the creation of a group of related virtual machines.

ISO files: CD/DVD image files are imported into this folder for virtual machine OS installation use.

Virtual machine files: These are configuration files for the virtual machines.

Virtual machine templates: Each template includes a fully installed and configured operating system with all the required applications and software stacks pre-installed. The templates are reusable for creating new virtual machines with exact image and configuration.

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

Oracle recommends a minimal of 10 GB for a storage repository and enough space should be allocated for a storage repository for virtual machines, templates, ISO files, and other virtual machine resources.

3.2.1 Storage efficiency with Dynamic Capacity

While Oracle VM allows for multiple storage repositories, only one physical disk is allowed in a repository. Therefore, a repository is limited by the size of the physical volume within it. With traditional storage, once the underlying physical disk in a repository is full, there is no easy way to expand the repository on the fly. A new repository is created with a bigger physical disk and then the data is migrated from the old to the new repository. On the other hand, if the physical disk is oversized, the spare disk space is wasted.

The SC Series Dynamic Capacity feature, also commonly known as thin provisioning, addresses this storage allocation and reservation challenge. SC volumes are allocated with a desired size but without consuming the actual storage upfront. The actual storage gets consumed only when the data is written to the volumes.

Therefore, it is perfectly fine to create and allocate a larger size volume to the repository. This allows room for it to grow over time without the fear of wasting storage.

When the volume is full, the volume can be expanded on the fly without taking any outages. This dramatically cuts down the time and impacts involved in moving files between repositories. An example of how to extend an Oracle VM storage repository online can be found in section 3.4.

Note: The SC Dynamic Capacity feature reduces storage consumption and improves utilization efficiency significantly.

3.2.2 Automated storage tiering for storage repository

The restriction of one disk per repository makes it difficult to optimize storage for all virtual machines. In traditional storage, each physical disk has a single performance characteristic. Therefore, all virtual machines in the same repository share the same performance trait regardless of what the applications require.

3.2.2.1 Benefits of SC volume and Data Progression

The SC Series array has a highly advanced and adaptable virtualization architecture which offers better performance and lowers storage utilization. Each virtualized volume can span across multiple physical disks of various different disk types and RAID levels and the balance between performance and cost is continuously re-evaluated and re-adjusted over time by the Data Progression feature.

The intelligent Data Progression algorithm tracks all block usage and automatically moves them to the optimum tiers and RAID levels based on their actual use. It guarantees all active writes go to the fastest disks at RAID 10 performance level. Blocks that have not been accessed for a certain period of time get redistributed to the lower cost disks at RAID 5 or RAID 6 performance levels. See Figure 8. As a result, virtual machines benefit from increased performance while keeping the cost low. Data Progression runs once a day to assess disk use and moves data to the storage tiers based on the storage profile applied to each volume.

Automated storage tiering and Data Progression are transparent to Oracle VM. No configuration within Oracle VM is necessary to take advantage of these advanced features.

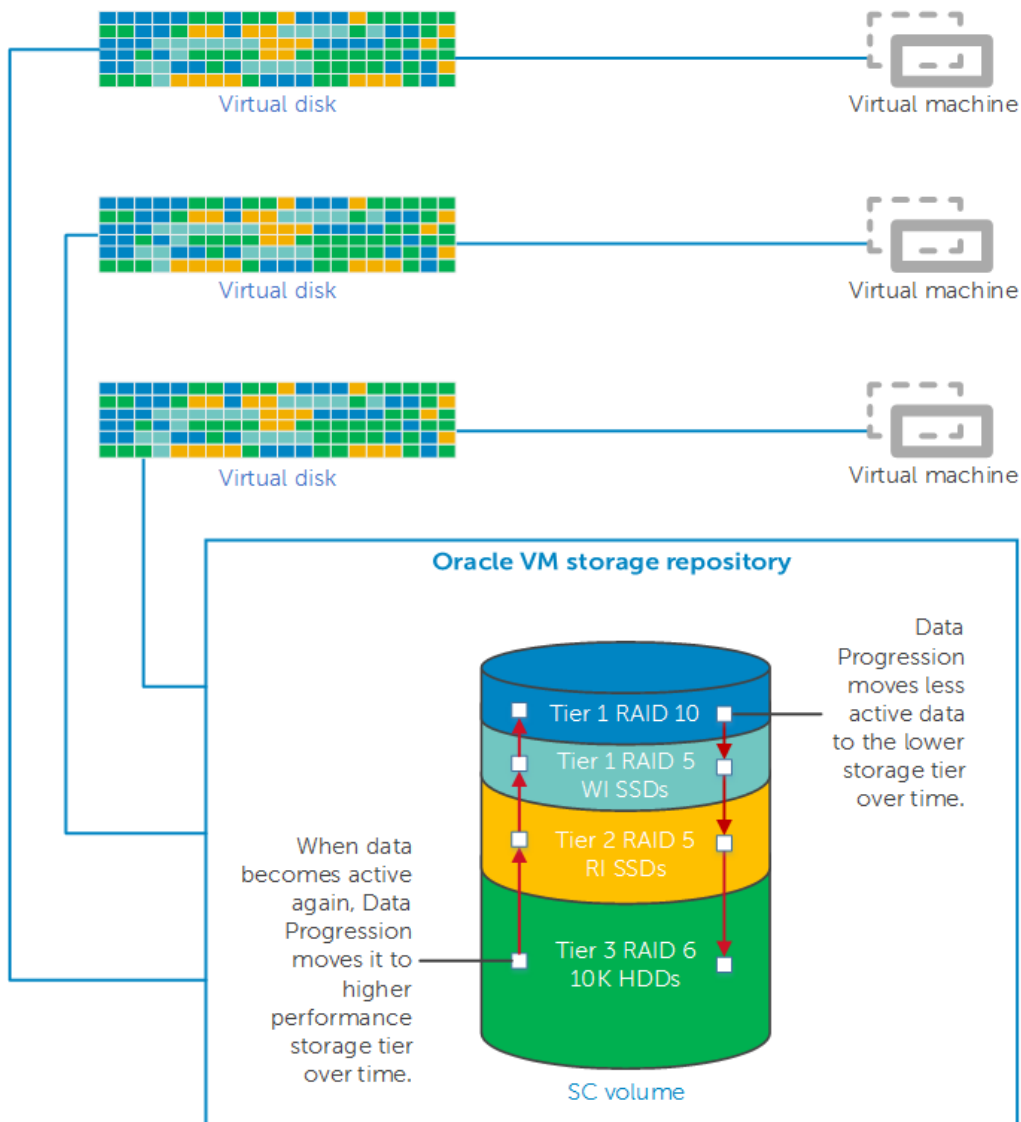


Figure 8 Automatic storage tiering and data progression in a Oracle VM storage repository

3.2.2.2 Choosing storage profiles

A storage profile is applied to each volume which determines how data is distributed and moved most effectively within a SC Series array. The following predefined profiles shown in Table 4 and Table 5 are available for standard storage types and flash-optimized storage types. Custom storage profiles can also be created if predefined profiles are not sufficient. Standard storage types include 7K, 10K, and 15K spinning drives. Flash storage types include write-intensive and read-intensive SSDs. A flash-optimized storage array has both flash storage and standard storage within the array.

Storage profiles also define where to hold Replays, which are block-level storage snapshots. The feature is discussed in section 4.1.

Table 4 Storage profiles for standard storage types

Profile name	Initial write tier	Tiers used and RAID levels	Progression
Recommended (All Tiers)	1	Writes: T1 RAID 10 Replays: RAID 5/RAID 6	Yes – all tiers
High Priority (Tier 1)	1	Writes: T1 RAID 10 Replays: T1 RAID 5/RAID 6	No
Medium Priority (Tier 2)	2	Writes: T2 RAID 10 Replays: T2 RAID 5/RAID 6	No
Low Priority (Tier 3)	3	Writes: T3 RAID 10 Replays: T3 RAID 5/RAID 6	No

Table 5 Storage profiles for flash-optimized storage types

Profile name	Initial write tier	Tiers used and RAID levels	Progression
Flash Only with Progression (Tier 1 to Tier 2)	1	Writes: T1 RAID 10 Replays: T2 RAID 5	Tier 1 to tier 2 only
Flash Optimized with Progression (Tier 1 to All Tiers)	1	Writes: T1 RAID 10 Replays: T2/T3 RAID 5/ RAID 6	Yes – all tiers
Write Intensive (Tier 1)	1	Writes: T1 RAID 10 Replays: T1 RAID 10	No
Low Priority (Tier 3)	3	Writes: T3 RAID 10 Replays: T3 RAID 5/RAID 6	No
Low Priority with Progression (Tier 3 to Tier 2)	3	Writes: T3 RAID 10 Replays: T3 RAID 5/RAID 6, or T2 RAID 5	Tier 3 to tier 2 only

Note: To maximize flexibility, performance, and cost-savings benefits, multiple types of drives should be installed in a SC Series array, and Storage Profiles that move data across all storage tiers should be chosen, such as **Recommended** or **Flash Optimized with Progression**.

It is recommended to have at least one Oracle VM storage repository that spans across all tiers. Additional repositories might be added with a different storage profile for specific needs. The following Table 6 shows some examples of different storage repository use cases.

Table 6 Multiple repository use cases

Repository	Storage profile	Used for	Description
Repo1	Recommended (All Tiers) or Flash Optimized with Progression	Virtual disks, virtual files	These files are frequently accessed by active virtual machines
Repo2	Low Priority or Low Priority with Progression	ISO files, assemblies, virtual machine templates	These files are usually large and do not get used often
Repo3	High Priority, Write Intensive, or Flash Only with Progression	High demanding applications/virtual machines	Isolate critical applications

Storage profiles can be changed at any time if the initial profile does not fit the need or the requirement has changed over time.

For more information on storage profiles, consult the *Dell Enterprise Manager Administrator's Guide*.

3.3 Creating Oracle VM storage repository with SC Series

This section walks through an example to create a storage repository with an SC Series array. All operations are performed in Dell Enterprise Manager and the Oracle VM Manager web console.

3.3.1 Creating an SC volume

In Dell Enterprise Manager, go to the **Storage** tab, right-click the **Volumes** node in the right pane, and select **Create Volume**. Enter the following information in the next window:

Name: Provide a user-friendly name for the volume. This name does not tie to anything on the Oracle VM Servers and will only be used in Dell Enterprise Manager.

Size: Provide the size of the volume to be allocated. Remember that it does not take up any disk space until the data is written to it.

Replay profile: A Replay is a storage-level snapshot and the Replay profile defines how often and what time Replays should be generated. This feature is discussed in detail in section 4. Select the default Replay profile if not sure.

Read/Write cache: When flash drives are used in the storage tiers, disabling write caching can improve performance. If only spinning drives are used, both read and write caching should be enabled.

Compression: Enabling compression can further reduce storage utilization. Not all data is eligible for compression. Only frozen pages by a Replay or inaccessible pages due to new data written over it are eligible for compression. For more information on this feature, consult the *Dell Enterprise Manager Administrator's Guide*.

Storage Profile: The Recommended or Flash Optimized with Progression storage profile is highly recommended, depending on the storage types installed in the SC Series array.

The screenshot shows the 'Create Volume' window in Dell Enterprise Manager. The window is divided into several sections:

- Name:** OVM-Disk1
- Size:** 500 GB
- Volume Folder:** A tree view showing the hierarchy: Volumes > Misc > OVM.
- Notes:** A text area for additional information.
- Replay Profiles:** A dropdown menu set to 'Daily'.
- Server:** A dropdown menu with 'Change' and 'Change' buttons.
- Read Cache:** A checkbox labeled 'Enabled'.
- Write Cache:** A checkbox labeled 'Enabled'.
- Compression:** A checkbox labeled 'Enabled'.
- Storage Profile:** A dropdown menu set to 'Recommended (All Tiers)'.

Figure 9 Create Volume window

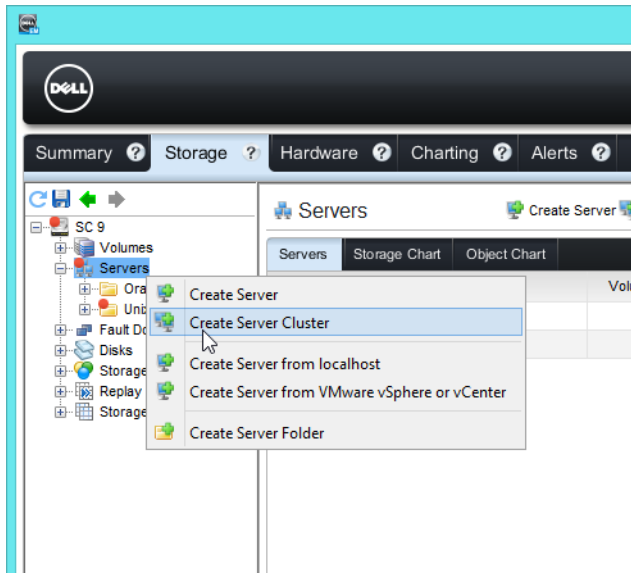
3.3.2 Creating server nodes and server cluster in Dell Enterprise Manager

Before volumes can be mapped and presented to the Oracle VM Servers, proper storage connectivity must be configured between the servers and storage array. Section 2 presents the various options and configurations available. Section 0 shows the steps to create a server node definition.

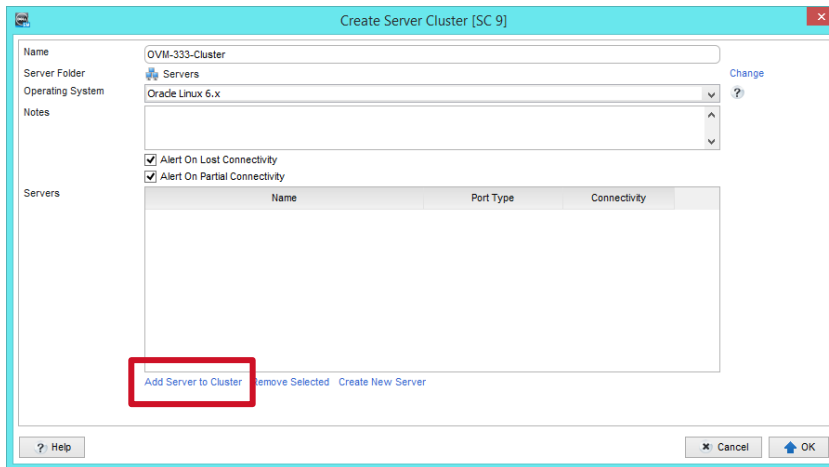
In order for virtual machines to migrate or fail over to any Oracle VM Servers in the cluster, map volumes to all servers in the same cluster so that the guest virtual machines can access the same storage on any host servers. To simplify the mapping operation, create a server cluster in Dell Enterprise Manager that consists of all the servers. Volumes can then be mapped to the server cluster in a single operation.

The following shows the steps to create a server cluster and map a volume to the server cluster in Dell Enterprise Manager:

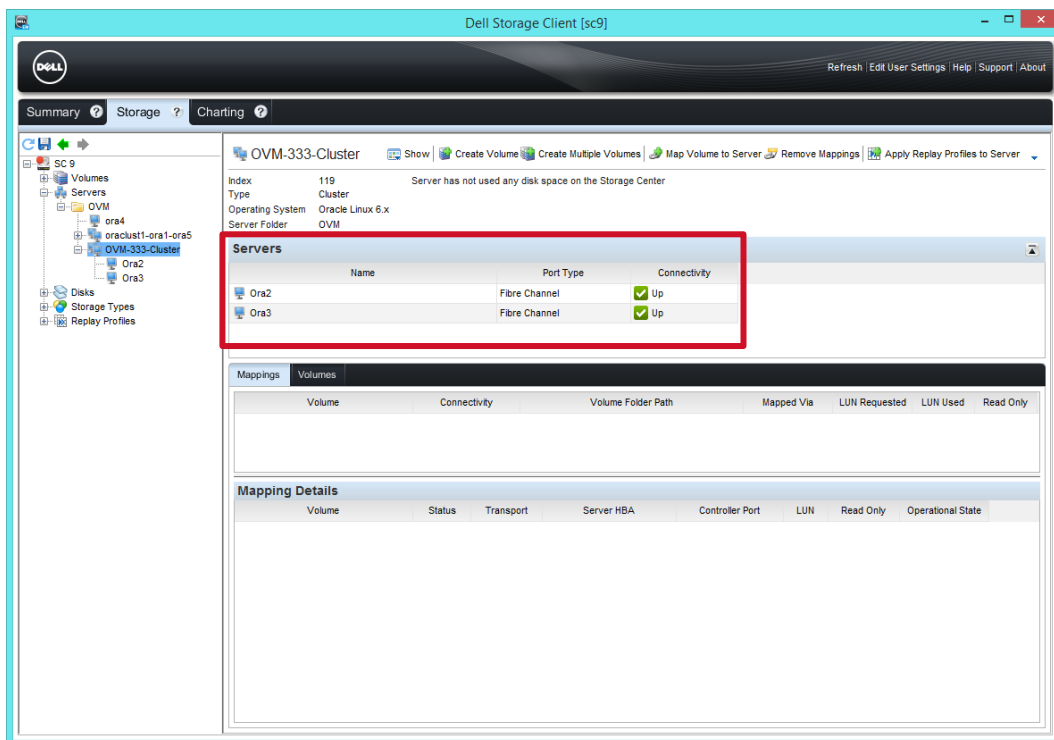
1. Under the **Storage** tab, select the **Servers** node in the left pane. Right-click **Servers** and select **Create Server Cluster**.



11. Fill out the **Name**, select the **Operating System**, and click the **Add Server to Cluster** link at the bottom.




12. Select and add the Oracle VM Servers to the cluster. In this example, the OVM-333-Cluster server cluster consists of Ora2 and Ora3 servers.



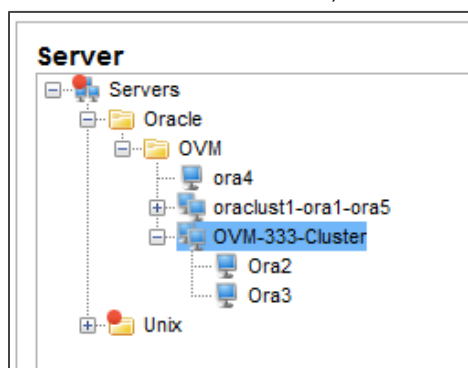
3.3.3 Mapping volumes to Oracle VM Servers

Volumes can now be mapped to any connected server or server cluster in Dell Enterprise Manager. Using a server cluster is more efficient when many servers are involved in the mapping process. The following steps demonstrate mapping a volume to a server cluster and making it available in Oracle VM.

1. Under the **Storage** tab, navigate to the **Volumes** node and the volume folder in the left pane. The right pane shows all volumes within the volume folder.
If mapping multiple volumes simultaneously, use Ctrl+click to select the volumes. Then, right-click to select **Map Volume to Server**.



13. In the next window, select the target server or server cluster.



14. In the next window, click the **Advanced** link to review all options. For multipathing configuration, set the **Maximum number of paths per Server** to **Unlimited**. If the server will be booting off this volume, the **Map volume using LUN 0** option must be checked.

The following volume and server will be mapped

Volume TestVol 1
Server OVM-333-Cluster

Select LUN

☒ Use next available LUN
☐ 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

Configure Multipathing

Maximum number of paths per Server

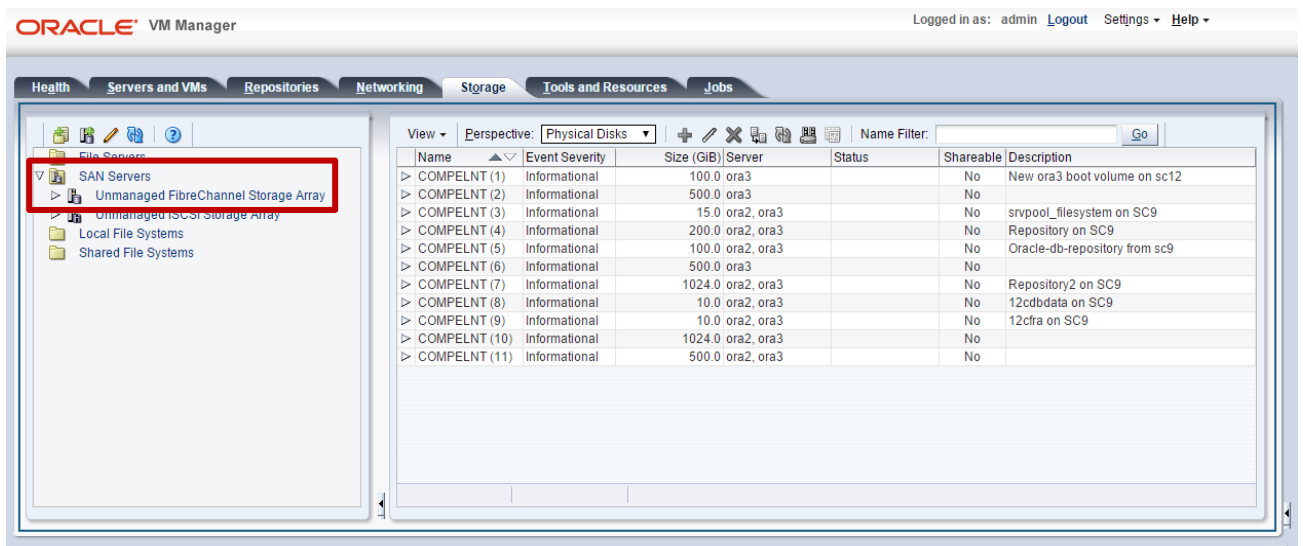
Configure Volume Use

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

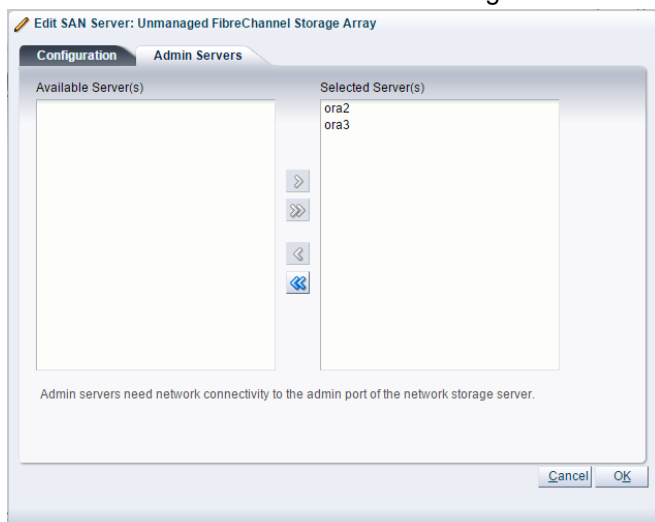
15. Click **Finish** to complete the mapping operation.

16. To scan for the new disk, in the Oracle VM Manager web console, 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.



17. For SC Series arrays, all Oracle VM Servers that have access to the storage should be added to the SAN Server admin server list. Right-click the SAN Server and edit the **Admin Servers** list.



18. Rescan and update storage on SAN Servers. In this example, an SC volume is mapped to the servers on Fibre Channel. Right-click **Unmanaged FibreChannel Storage Array** and select **Refresh**. Once the refresh is done, the new volume shows up under the **Physical Disks** perspective. See Figure 10.

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.

- Click the arrow on the leftmost column to expand the disk entry to reveal the details of the disk. The **Page83 ID** is the WWN of the volume which is also part of the SC volume device ID.

The image consists of two screenshots. The top screenshot is from the Oracle VM Manager interface, showing a table of physical disks. The bottom screenshot is from the Dell Storage Client interface, showing details for a specific volume.

Oracle VM Manager Screenshot:

Name	Event Severity	Size (GiB)	Server	Status	Shareable	Description
COMPELNT (1)	Informational	100.0	ora3		No	New ora3 boot volume on sc12
COMPELNT (2)	Informational	500.0	ora3		No	
COMPELNT (3)	Informational	15.0	ora2, ora3		No	srvpool_filesystem on SC9
COMPELNT (4)	Informational	200.0	ora2, ora3		No	Repository on SC9
COMPELNT (5)	Informational	100.0	ora2, ora3		No	Oracle-db-repository from sc9
COMPELNT (6)	Informational	500.0	ora3		No	
COMPELNT (7)	Informational	1024.0	ora2, ora3		No	Repository2 on SC9
COMPELNT (8)	Informational	10.0	ora2, ora3		No	12cddbdata on SC9
COMPELNT (9)	Informational	10.0	ora2, ora3		No	12cfra on SC9
COMPELNT (10)	Informational	1024.0	ora2, ora3		No	
COMPELNT (11)	Informational	500.0	ora2, ora3		No	

Below the table, the details for COMPELNT (11) are shown:

- Name: COMPELNT (11)
- User Friendly Name: 36000d3100000650000000000000201a
- San Server: Unmanaged FibreChannel Storage Array
- Thin Provision: Yes
- Type: LUN
- ID: 0004b00001800003c0ddfb2073c740
- Page83 ID: 36000d3100000650000000000000201a

Dell Storage Client Screenshot:

The Dell Storage Client interface shows the following details for TestVol 1:

- Index: 8652
- Configured Size: 500 GB
- Volume Folder: OVM-3.3.3
- Volume Active On Controller: 102

The **General** tab shows:

- Serial Number: 00000065-0000201a
- Device ID: 6000d31000006500000000000000201a
- Replay Profile List: Daily
- Storage Profile: Recommended (All Tiers)
- Storage Type: Assigned - Redundant - 2 MB
- Disk Folder: Assigned
- Compression Enabled: No
- Host Cache Enabled: No
- Created On: 1/15/16 12:22:33 PM
- Modified On: 1/15/16 12:22:33 PM

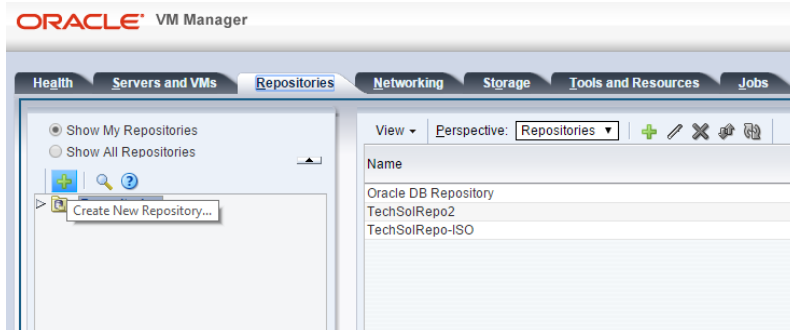
A red arrow points from the Page83 ID in the Oracle VM Manager screenshot to the Device ID in the Dell Storage Client screenshot, indicating that the Page83 ID is the WWN of the volume which is also part of the SC volume device ID.

Figure 10 Identifying an SC volume in Oracle VM Manager

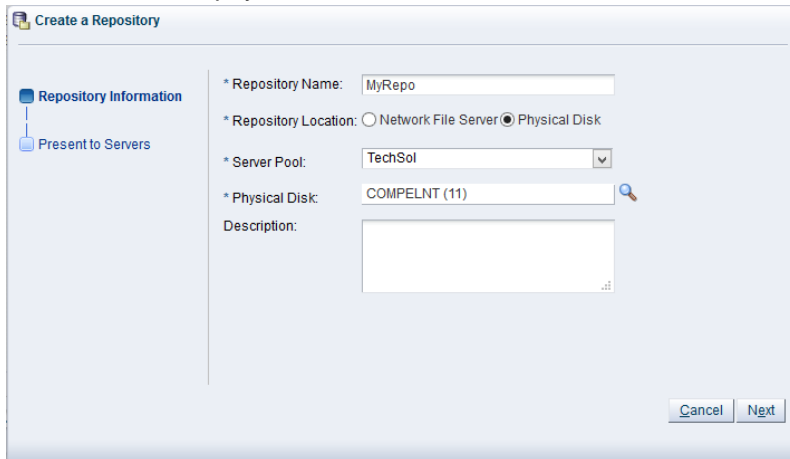
3.3.4 Creating a new repository and present it to a server pool

The following steps demonstrate using the previously mapped volume to create a new storage repository.

1. In Oracle VM Manager web console, under the **Repositories** tab, click the **+** icon in the left pane to create a new repository.



19. Choose **Physical Disk** for **Repository Location**. Click the magnifying lens icon to see a list of available physical disks.



20. Select **Unmanaged FibreChannel Storage Array** and the volume mapped previously.

Create a Repository: Select Physical Disk

SAN Server: **Unmanaged FibreChannel Storage Array**
 Volume Group:
 Name Filter:

Name	Size (GiB)	User Friendly Name
COMPELNT (1)	100.0	36000d3100002b9000000000000006393a
COMPELNT (6)	500.0	36000d3100000650000000000000002019
COMPELNT (2)	500.0	36000d3100000650000000000000002018
COMPELNT (10)	1024.0	36000d3100000650000000000000001f4c
COMPELNT (11)	500.0	36000d310000065000000000000000201a

21. During its creation, the storage repository requires a server pool to present to. Therefore, a server pool must be configured and available for this to continue. Consult the Oracle VM administration guide for information and instructions on creating a server pool.
22. The new repository shows up in the Repositories perspective. Notice the File System and the overall usage information. A storage repository is essentially an Oracle Cluster File System version 2 (OCFS2) file system created on the disk. When the repository is presented to the server pool, each server in the pool mounts the OCFS2 file system on the server.

ORACLE VM Manager Logged in as: admin [Logout](#) [Settings](#) [Help](#)

Health Servers and VMs **Repositories** Networking Storage Tools and Resources Jobs

Show My Repositories
Show All Repositories

Repositories

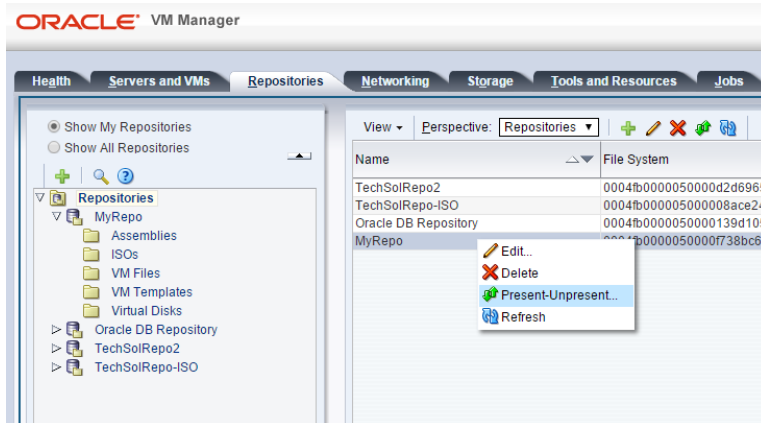
- MyRepo
 - Assemblies
 - ISOs
 - VM Files
 - VM Templates
 - Virtual Disks
- Oracle DB Repository
- TechSolRepo2
- TechSolRepo-ISO

View Perspective: **Repositories**

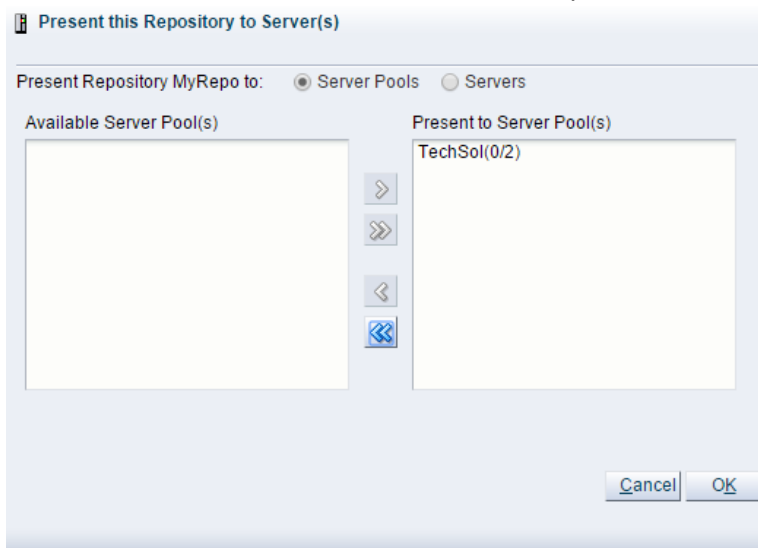
Name	File System	File System Size (GiB)		
		Free	Used	Total
TechSolRepo2	0004fb0000050000d2d69657212a452c [fs_TechSolRepo2]	529.64	463.64	993.28
TechSolRepo-ISO	0004fb000005000008ace24432088547 [fs_TechSolRepo]	109.04	90.96	200.0
Oracle DB Repository	0004fb0000050000139d1054c145da4e	92.14	7.86	100.0
MyRepo	0004fb0000050000738bc6136c71805	494.47	5.53	500.0

Columns Hidden 1

23. For Oracle VM to be able to use the storage repository, it must be presented to a server pool. Right-click the repository in the **Repositories** perspective and select the **Present-Unpresent** option.



24. Choose from the list of available server pools or servers.



The repository is now available for use. The OCFS2 file system can be seen mounted on the Oracle VM Servers.

25. ssh to one of the Oracle VM Servers and execute the `df` command to list the mounted storage repository. In this example, MyRepo repository is mounted on `/OVS/Repositories/0004fb0000030000ea56d4a32481e0b3`.

```
[root@ora2 nodes]# df -m
Filesystem                                1M-blocks  Used Available Use% Mounted on
/dev/sda2                                50269      2874    44819    7% /
tmpfs                                     995         0     995      0% /dev/shm
/dev/sda1                                477       125     323    28% /boot
none                                     995         7     988     1% /var/lib/xenstored
/dev/mapper/36000d3100000650000000000000017dc 15360      263    15098     2% /poolfsmnt/0004fb0000050000aa64db01ff45b9c8
/dev/mapper/36000d3100000650000000000000017df 204800    93145   111655   46% /OVS/Repositories/0004fb00000300006cdf01daced67181
/dev/mapper/36000d310000065000000000000001f08 1017119  489100   528019   49% /OVS/Repositories/0004fb0000030000a7e754fc1745a7c8
/dev/mapper/36000d31000006500000000000000150d 103100     8053    94047    8% /OVS/Repositories/0004fb0000030000638e714565e59956
/dev/mapper/36000d31000006500000000000000201a 512000    5663   506337    2% /OVS/Repositories/0004fb0000030000ea56d4a32481e0b3
[root@ora2 nodes]#
```

3.4 Extending Oracle VM storage repository

Since each Oracle VM storage repository can only have one volume, when the capacity of the repository reaches the limit, the administrator would either have to create another repository and migrate existing files to the new repository or extend the capacity of the existing repository by growing the size of the volume underneath it.

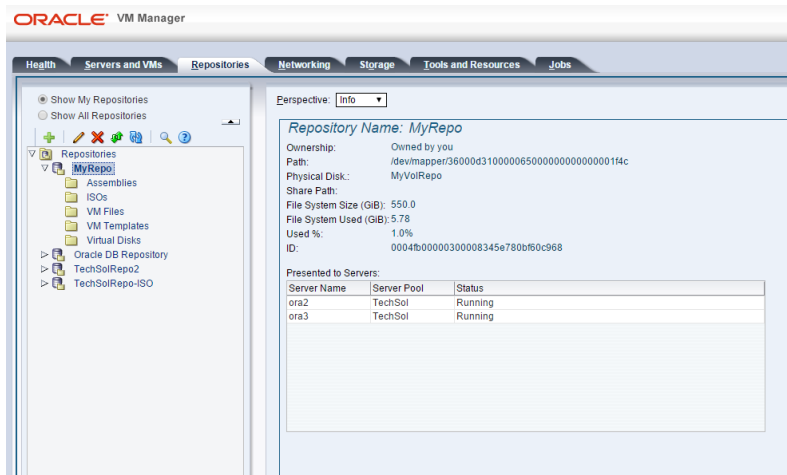
The operation of growing the volume size can be done online in Dell Enterprise Manager. Extending an Oracle VM repository can also be done online without incurring any outages in version 3.3.

Use the following steps to extend an Oracle VM storage repository online:

1. Identify the repository file system and volume.

In the Oracle VM console, go to the **Repositories** tab and select the Repository in the left pane. The **Path** shows the OCFS2 file system information.

The volume's device id can be extracted from the **Path** if the device is multipathed and does not use friendly name alias. Otherwise, additional **Physical Disk** information including the device id can be found in the **Storage** tab. The device id can be used to correlate the volume in Dell Enterprise Manager.



26. Log in to one of the Oracle VM Servers (optional)

This is another way to display the information in step 1. To confirm the repository file system on an Oracle VM Server, log in through `ssh` and show the information with native Linux commands.

- Use `fdisk` to show the current physical disk size:

```
# fdisk -l /dev/mapper/36000d310000065000000000000001f4c
```

```
Disk /dev/mapper/36000d310000065000000000000001f4c: 590.6 GB,  
590558003200 bytes  
255 heads, 63 sectors/track, 71797 cylinders  
Units = cylinders of 16065 * 512 = 8225280 bytes  
Sector size (logical/physical): 512 bytes / 4096 bytes  
I/O size (minimum/optimal): 2097152 bytes / 2097152 bytes  
Disk identifier: 0x02020202
```

- Use `df` to show the repository file system size:

```
# df -k
Filesystem                                1K-blocks      Used
    Available Use% Mounted on
/dev/mapper/36000d3100002b9000000000000006393ap2 25396260  2055012
    22030344   9% /
tmpfs                                           869448         0
    869448   0% /dev/shm
/dev/mapper/36000d3100002b9000000000000006393ap1 1019208  151892
    814708  16% /boot
none                                           869448        168
    869280   1% /var/lib/xenstored
/dev/mapper/36000d310000065000000000000000017dc 15728640  268908
    15459732   2% /poolfsmnt/0004fb0000050000aa64db01ff45b9c8
/dev/mapper/36000d31000006500000000000000001f4c 576716800  6063104
    570653696   2% /OVS/Repositories/0004fb00000300008345e780bf60c968
```

- Use `multipath` to show all disk paths and the device id of the physical disk:

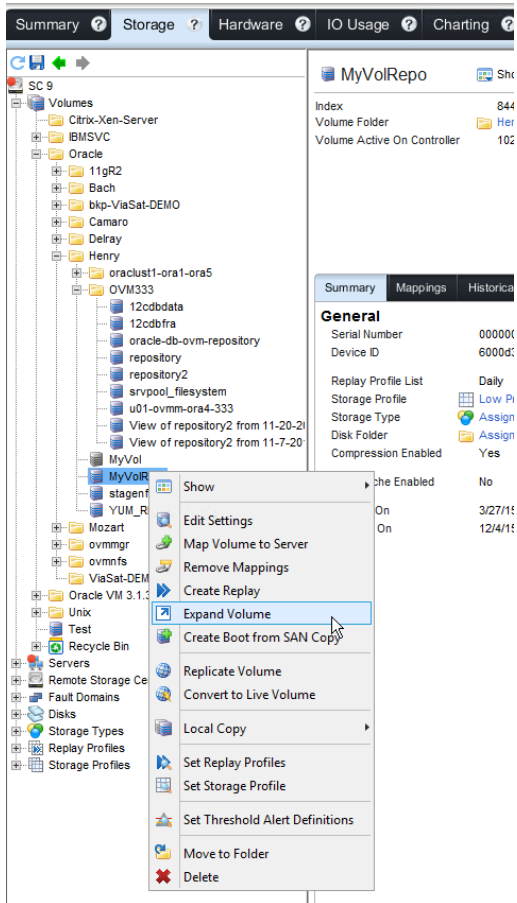
```
# multipath -ll 36000d31000006500000000000000001f4c
36000d31000006500000000000000001f4c dm-14 COMPELNT,Compellent Vol
size=550G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
    |- 2:0:1:7 sdi 8:128 active ready running
    |- 2:0:2:7 sdn 8:208 active ready running
    |- 3:0:1:7 sdad 65:208 active ready running
    `- 3:0:2:7 sdai 66:32 active ready running
```

- Use `scsi_id` to extract the device id:

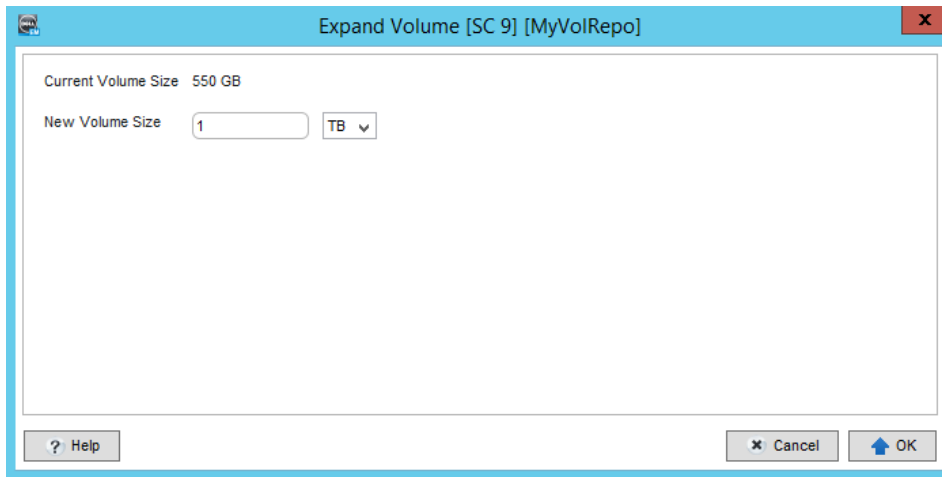
```
# scsi_id --page=0x83 --whitelisted \
    --device=/dev/mapper/36000d31000006500000000000000001f4c
```

27. Extend the volume in Dell Enterprise Manager.

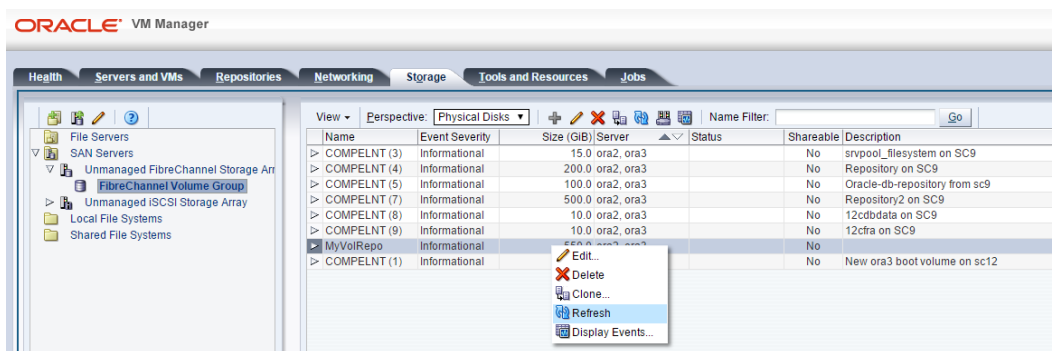
- a. In Dell Enterprise Manager, locate the volume corresponding to the repository and select **Expand Volume** from the right-click menu.



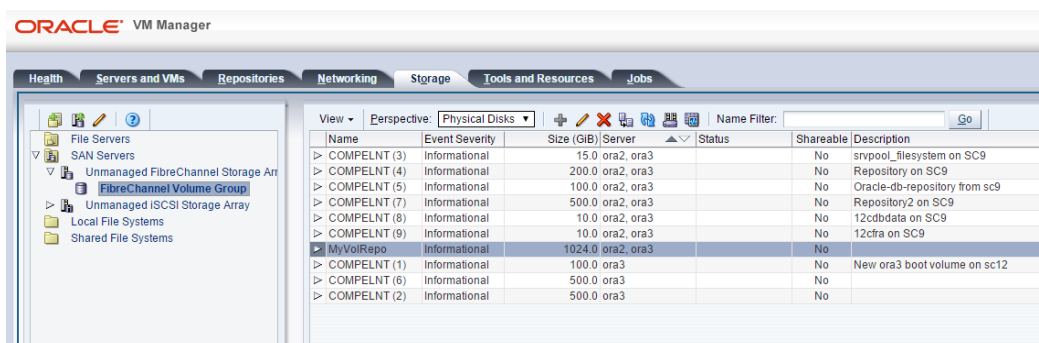
- b. Enter the new volume size in the pop-up window. In this example, the original size of the volume is 550 GB and the new size is 1 TB.



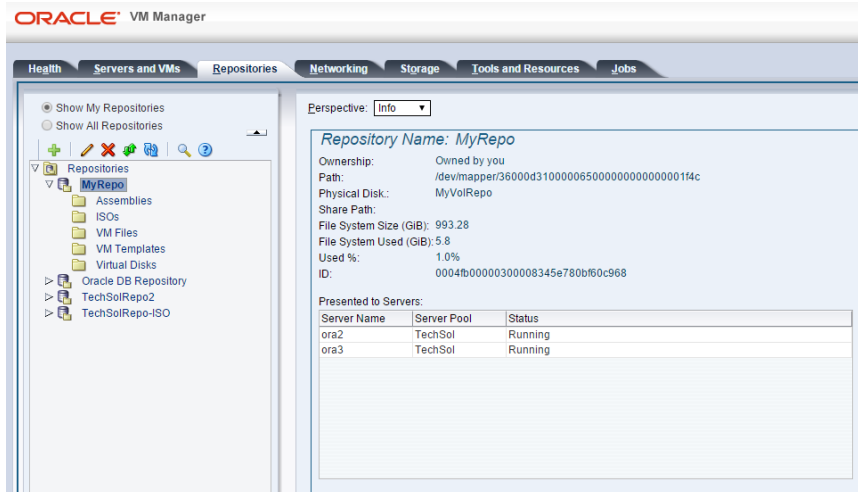
28. Resize the repository file systems in Oracle VM console.
Under the **Storage** tab, locate the volume and select **Refresh** from the right-click menu. This tells Oracle VM Servers to reread the physical disk information and extend the file system on it.



The new size is now reflected in the list.



29. Verify the new size of the repository in the **Repositories** tab.



30. Verify repository size through `ssh` (optional).

Log in to the Oracle VM Servers using `ssh` and run the native Linux commands to confirm the new size.

```
# multipath -ll 36000d31000006500000000000000001f4c
```

```
36000d31000006500000000000000001f4c dm-14 COMPELNT,Compellent Vol
size=1.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
   |- 2:0:1:7 sdi 8:128 active ready running
   |- 2:0:2:7 sdn 8:208 active ready running
   |- 3:0:1:7 sdad 65:208 active ready running
   `-- 3:0:2:7 sdai 66:32 active ready running
```

```
# df -k
```

Filesystem	1K-blocks	Used
Available Use% Mounted on		
/dev/mapper/36000d3100002b9000000000000006393ap2	25396260	2055012
22030344 9% /		
tmpfs	869448	0
869448 0% /dev/shm		
/dev/mapper/36000d3100002b9000000000000006393ap1	1019208	151892
814708 16% /boot		
none	869448	168
869280 1% /var/lib/xenstored		
/dev/mapper/36000d3100000650000000000000017dc	15728640	268908
15459732 2% /poolfsmnt/0004fb0000050000aa64db01ff45b9c8		
/dev/mapper/36000d310000065000000000000001f4c	1041529856	6077440
1035452416 1% /OVS/Repositories/0004fb00000300008345e780bf60c968		

3.5 Oracle VM Manager storage consideration

The Oracle VM Manager is a web-based management console that manages and monitors the Oracle VM Servers and virtual machines. The Oracle VM Manager applications consist of an Oracle WebLogic Server and an Oracle database or MySQL database for the management repository. The MySQL database has become the default bundled database starting with Oracle VM Manager 3.2.

To simplify deployment, it is recommended to install and use the bundled database for the management repository. Oracle fully supports the use of the bundled MySQL database in a production environment.

The Oracle VM Manager applications are typically installed on a separate physical server that supports Oracle Linux 5.5 and above. Here are some recommendations and guidelines:

- Install the Oracle VM Manager applications on one or more dedicated SC volumes and take advantage of auto-tiering, thin provisioning, storage snapshot, and other advanced features.
- A single file system may be sufficient for all applications and the management repository.
- To increase performance and improve storage management, create the file system on Linux Logical Volume Manager (LVM). LVM makes use of multiple disks and spreads the file system across them. This increases the number of available paths and I/O queues to the SAN storage.
- A minimum of two SC volumes is recommended when the SC Series array is a dual-controller system. This maximizes the utilization of both controllers because each volume is assigned to a different primary controller.
- Apply the same best practices and guidelines discussed in this paper and the document, [*Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 6x Best Practices*](#).

4 Protecting Oracle VM with storage snapshots

4.1 Data Instant Replay

Data Instant Replay, also known as Replay, is a block-level storage snapshot technology available on the SC Series array. This document uses Replay and snapshot interchangeably. A Replay is a point-in-time copy of a volume. Taking a Replay is fast and incurs no overhead on the servers. Multiple Replays can be taken over time and used as recovery points in the event of data corruption or loss due to events such as hardware failure or human errors.

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

Using Replays with Oracle VM Server storage repositories and VM Manager database offers many benefits:

- Space-efficient, time-based snapshots can be created at any time.
- Replay events can be scheduled using a Replay Profile in Dell Enterprise Manager, or manually created when needs arise.
- For a flash-optimized storage type where multiple types of flash storage and spinning drives are used in the same array, Replays can occur on demand by Storage Center to automatically move data from tier 1 to free up space on the flash storage.
- Data Progression and Replays work together seamlessly to move frozen data to the storage tier specified by the storage profile. For example, Replays can move from RAID 10 on tier 1 to RAID 5 in tier 3 to reduce storage consumption.
- A replay can be created in only a few seconds.
- Multiple volumes can be snapped together using the Consistent Replay profile. This is particularly important for databases that reside on more than one volume.
- Replays provide extra protection to the conventional backup/restore mechanism.
- Recovery from a Replay is fast and can minimize system downtime.
- Replays also work with remote SC replication for offsite storage and disaster recovery.

4.2 Taking snapshots of Oracle VM repositories and Oracle VM Manager database

The purpose of taking snapshots is to preserve a point-in-time copy of the data. The process of taking snapshots can be automated in Dell Enterprise Manager by defining and applying a Replay profile to each volume.

A Replay profile defines which volumes should be included, when the Replay should be taken, how often is they are scheduled, and the creation method. Each Replay also has an expiration date to facilitate automatic cleanup.

Note: It is recommended to have at least a daily Replay for Replay Data Progression to work properly.

4.2.1 Create a new Replay profile

1. In Dell Enterprise Manager under the **Storage** tab, right-click the **Replay Profile** node in the left pane and select **Create Replay Profile**.
2. One or more rules (schedules) can be defined. In the example, the schedule runs every hour, all day, every day. Any Replays older than two weeks are to be removed.

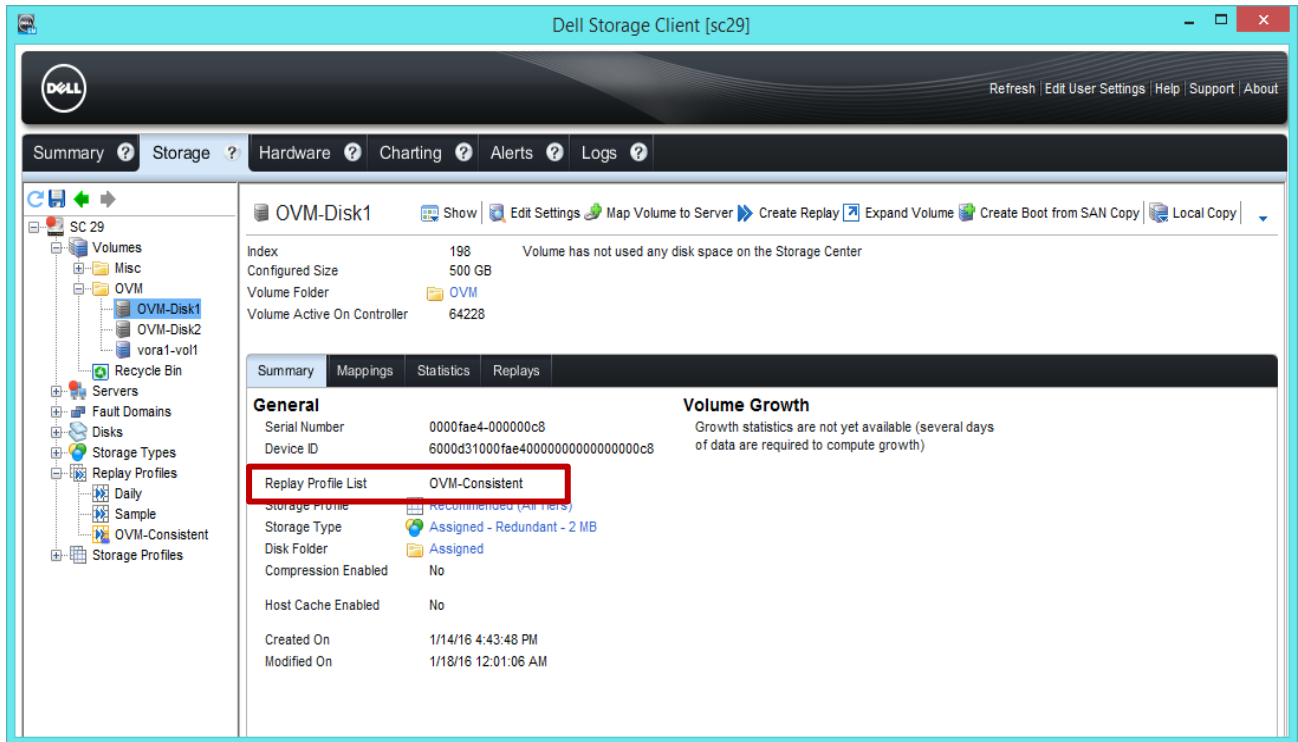
To take consistent snapshots for all volumes assigned to this profile, make sure the **Replay Creation Method** is set to **Consistent**.

The screenshot shows the 'Create Replay Profile' dialog box. It has a title bar with a close button. The main area is divided into sections: 'Name' (OVM-Consistent), 'Notes' (empty text area), and 'Rules'. The 'Rules' section contains a table with the following data:

Name	Expiration	Frequency	Time
Daily	2 weeks	Daily	Interval 1 hour Start Time 12:00 AM End Time 12:00 AM

Below the table are links: '+ Add Rule', 'Edit Rule', and '- Remove Rule'. The 'Replay Creation Method' is set to 'Consistent'. There are two checked checkboxes: 'Timeout Replay creation after 20 Seconds' and 'Expire incomplete Replay sets if Replay creation timeout is reached'. At the bottom are buttons for '? Help', 'Cancel', and 'OK'.

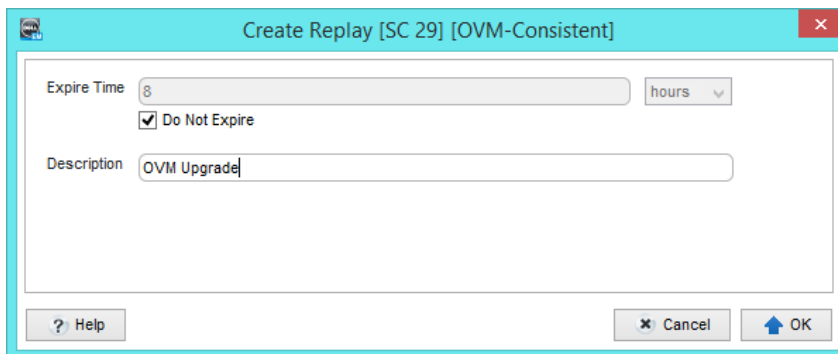
- To apply the profile to volumes, right-click the **Relay Profile** in the left pane and select **Apply to Volumes**. Select the volumes in the next pop-up window. Go to the summary page of a volume to see the Replay profiles applied to that volume.



4.2.2 Taking an on-demand snapshot

On-demand snapshots can be created outside of the schedule. This is useful for saving a quick restore point before making any planned changes.

Under the **Storage** tab, right-click the **Replay Profile** and select **Create Replay**. Specify an expiration time or make it never expire.



4.2.3 Listing all available replays

To see a list of Replays for a specific volume, click the volume in the left pane and the **Replays** tab in the right pane.

Freeze Time	Expire Time	Size	% of Actual	Description	Space Recovery	Consistent	Source	Create Volume	Replay Profile	Replay Profile Rule	Writes Held Duration
Active		0 MB	0%		No	No	System	OVM-Disk1			
1/18/16 9:00:01 AM	2/1/16 9:00:01 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.001 seconds
1/18/16 8:59:23 AM	1/18/16 4:59:23 PM	0 MB	0%	OVM Upgrade	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.003 seconds
1/18/16 8:00:00 AM	2/1/16 8:00:00 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.003 seconds
1/18/16 7:00:01 AM	2/1/16 7:00:01 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.004 seconds
1/18/16 6:00:00 AM	2/1/16 6:00:00 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.003 seconds
1/18/16 5:00:01 AM	2/1/16 5:00:01 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.003 seconds
1/18/16 4:00:00 AM	2/1/16 4:00:00 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.004 seconds
1/18/16 3:00:01 AM	2/1/16 3:00:01 AM	0 MB	0%	Daily	No	Yes	Consistenc...	OVM-Disk1	OVM-Consistent	Daily	0.001 seconds

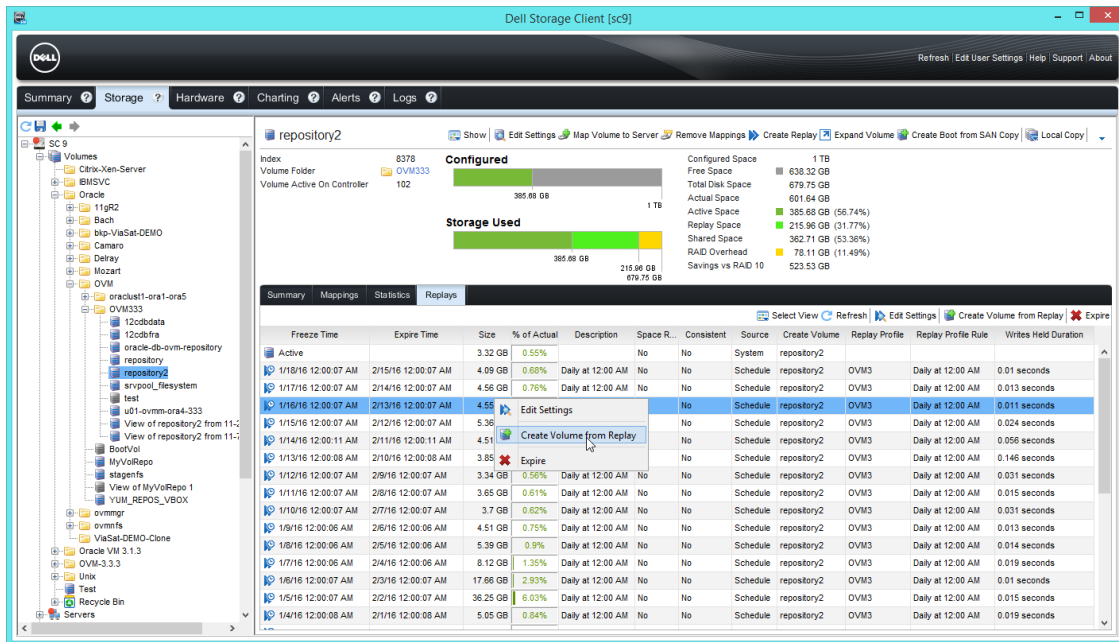
4.3 Restoring and recovering Oracle VM files from a snapshot

Once Replays are created, they can be used to restore loss or corrupted data. For Oracle VM, a virtual machine's configuration files and its virtual disks are stored in a storage repository. With a Replay, a view volume can be created that represents a copy of the repository from a specific point in time. If there are no changes to the data on the view volume, no storage is consumed. Only changes to the data on the view volume will actually consume space. The view volume is then mounted on the Oracle VM Server alongside the original repository. To restore a specific virtual machine or virtual disks, simply copy the files from the restored repository to the original repository.

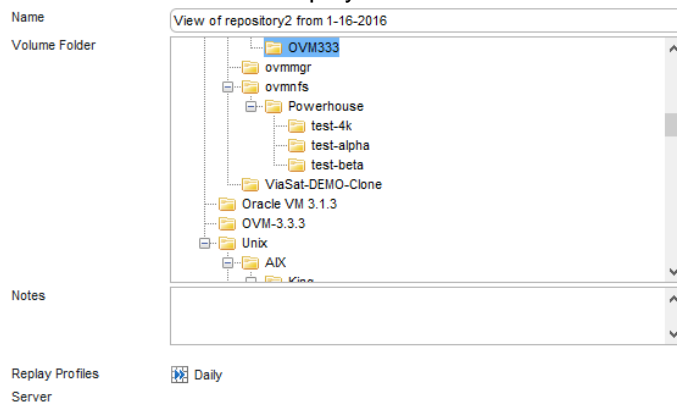
The following sections demonstrate the steps to restore a deleted virtual disk of a virtual machine from a replay. The same steps can be applied to other files in a repository.

4.3.1 Creating a view volume from a Replay

1. In Dell Enterprise Manager, navigate to the volume in the left pane. Under the **Replays** section in the right pane, right-click the Replay for a specific time and select **Create Volume from Replay**.

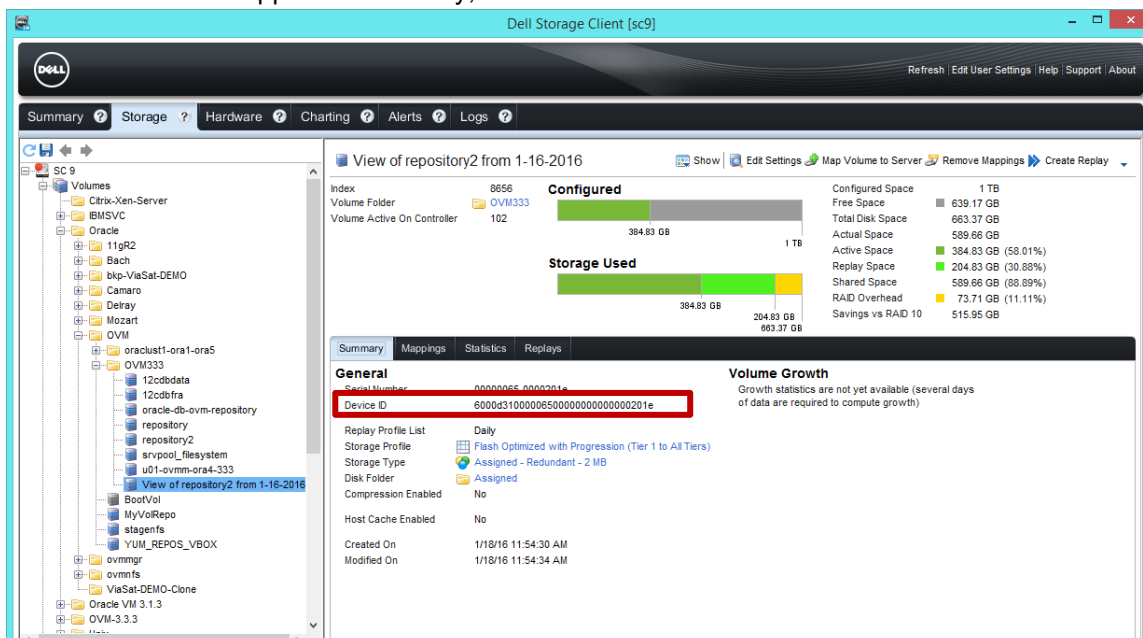


2. Provide a name for the view volume in the next window. It is a best practice to incorporate the date from which the Replay was taken.

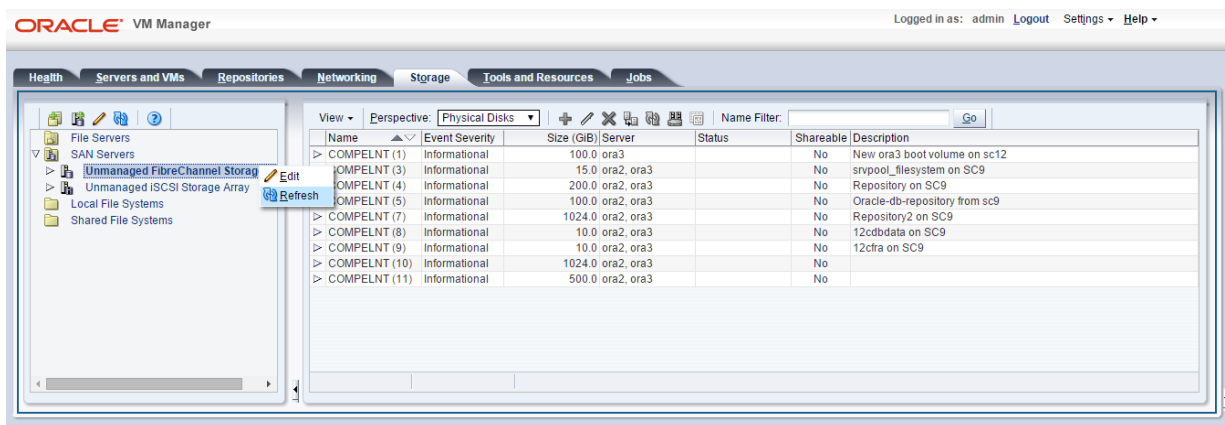


4.3.2 Mapping the view volume to the Oracle VM Server

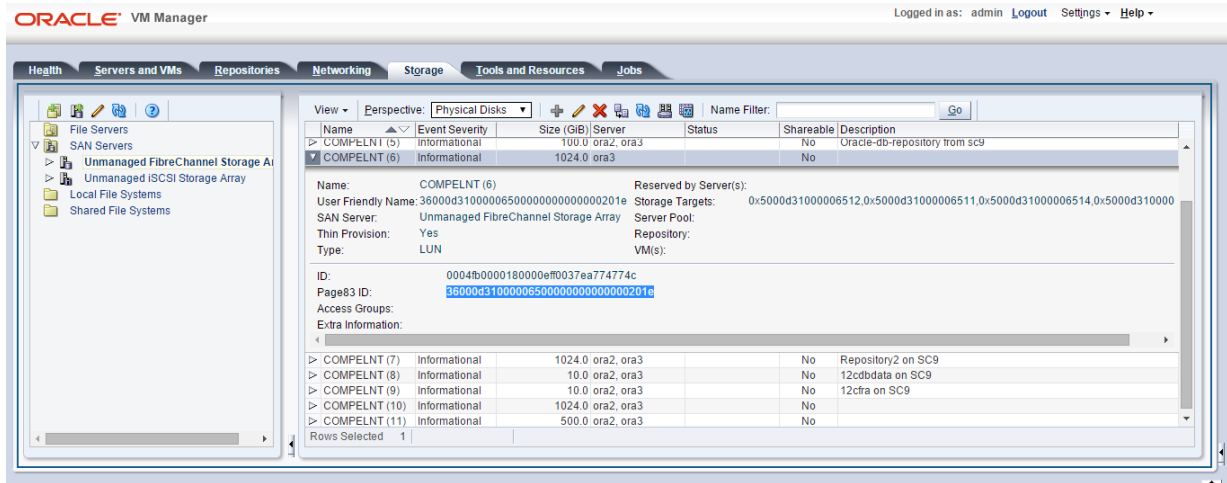
1. Right-click the view volume and select **Map Volume to Server**. Select one of the Oracle VM Servers from the list. It is sufficient to map the volume to one server in the server pool. In this example, the view volume is mapped to server ora3.
31. Once it is mapped successfully, take note of the volume's **Device ID**.



32. Scan for new volumes in Oracle VM Manager web console. Under the **Storage** tab, right-click the SAN Servers and select **Refresh**.



33. Identify the new volume after refresh is complete. In this example, **COMPELNT (6)** shows up after refresh and its **Page83 ID** matches the volume device ID.



4.3.3 Mounting the view volume on a Oracle VM Server

1. Change the volume label for OCFS2. Since the view volume essentially is a clone of the original volume, both original and view volumes contain the same OCFS2 label. In order for Oracle VM to recognize and mount the view volume, the label must be changed.
 - a. `ssh` to the Oracle VM Server that has the view volume. Use the `tuneufs.ocfs2` command to change the label. The command requires the volume's multipath device as an argument. Find the corresponding multipath device in the `/dev/mapper` directory.

```
# 11 /dev/mapper
```

```
total 0
lrwxrwxrwx 1 root root          7 Jan 18 12:01
    36000d31000006500000000000000017dc -> ../dm-6
lrwxrwxrwx 1 root root          8 Jan 18 12:01
    36000d31000006500000000000000017df -> ../dm-14
lrwxrwxrwx 1 root root          8 Jan 18 12:01
    36000d3100000650000000000000001f08 -> ../dm-15
lrwxrwxrwx 1 root root          7 Jan 18 12:01
    36000d3100000650000000000000001f0b -> ../dm-2
lrwxrwxrwx 1 root root          7 Jan 18 12:01
    36000d3100000650000000000000001f0c -> ../dm-1
lrwxrwxrwx 1 root root          7 Jan 18 12:01
    36000d3100000650000000000000001f0d -> ../dm-9
lrwxrwxrwx 1 root root          7 Jan 18 12:01
    36000d3100000650000000000000001f4c -> ../dm-7
lrwxrwxrwx 1 root root          8 Jan 18 12:01
    36000d310000065000000000000000201a -> ../dm-18
lrwxrwxrwx 1 root root          7 Jan 18 12:01
    36000d310000065000000000000000201b -> ../dm-5
```

```

lrwxrwxrwx 1 root root      8 Jan 18 12:01
36000d310000065000000000000000201e -> ../dm-13
lrwxrwxrwx 1 root root      7 Jan 18 12:01
36000d3100002b90000000000000006393a -> ../dm-8
lrwxrwxrwx 1 root root      8 Jan 18 12:01
36000d3100002b90000000000000006393ap1 -> ../dm-10
lrwxrwxrwx 1 root root      8 Jan 18 12:01
36000d3100002b90000000000000006393ap2 -> ../dm-11
lrwxrwxrwx 1 root root      8 Jan 18 12:01
36000d3100002b90000000000000006393ap3 -> ../dm-12
crw-rw---- 1 root root 10, 236 Jan  7 09:18 control
lrwxrwxrwx 1 root root      8 Jan 18 12:01 SATA_MZ-5EA1000-
OD3_S0SEN6AD401998 -> ../dm-17
lrwxrwxrwx 1 root root      7 Jan 18 12:01 SATA_MZ-5EA1000-
OD3_S0SEN6AD401998p1 -> ../dm-0
lrwxrwxrwx 1 root root      7 Jan 18 12:01 SATA_MZ-5EA1000-
OD3_S0SEN6AD401998p2 -> ../dm-3
lrwxrwxrwx 1 root root      7 Jan 18 12:01 SATA_MZ-5EA1000-
OD3_S0SEN6AD401998p3 -> ../dm-4

```

b. Update the label.

```

# tuneefs.ocfs2 --cloned-volume
/dev/mapper/36000d310000065000000000000000201e

```

Updating the UUID and label on cloned volume

```
"/dev/mapper/36000d310000065000000000000000201e".
```

DANGER: THIS WILL MODIFY THE UUID WITHOUT ACCESSING THE CLUSTER
SOFTWARE. YOU MUST BE ABSOLUTELY SURE THAT NO OTHER NODE IS USING
THIS FILESYSTEM BEFORE MODIFYING ITS UUID.

Update the UUID and label? **yes**

34. Mount the view volume.

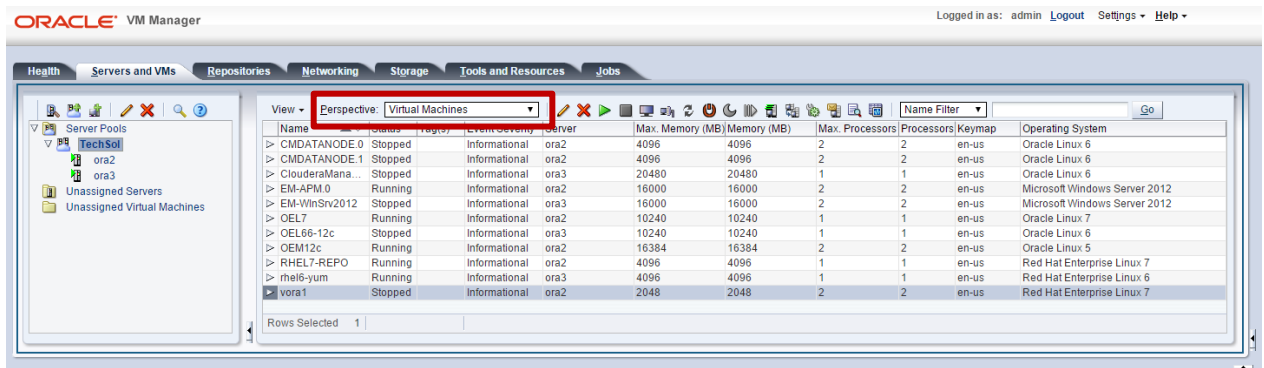
```

# mkdir /mnt/recover
# mount -t ocfs2 -o rw,heartbeat=global
/dev/mapper/36000d310000065000000000000000201e \ /mnt/recover

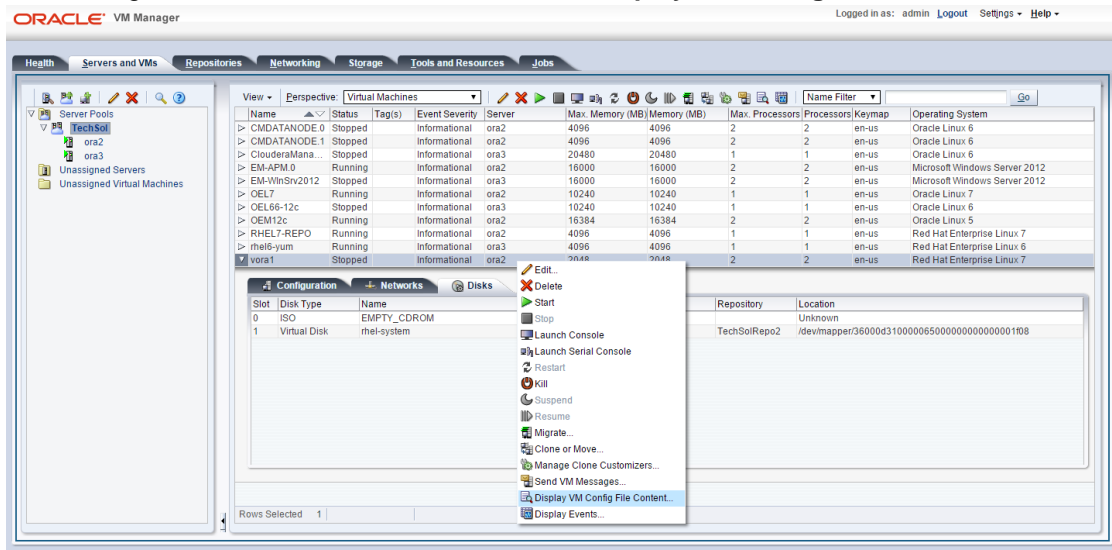
```

4.3.4 Performing recovery

1. First, identify the virtual disk file that needs to be restored and replaced in Oracle VM Manager web console. In this example, the data in virtual machine **vora1** has been corrupted.
To list the virtual disks used by vora1 virtual machine, in Oracle VM Manager web console, under the **Servers and VMs** tab, navigate to the **Server Pools** folder and change the **Perspective** to **Virtual Machines**.



35. Right-click the virtual machine and select **Display VM Config File Content**.



36. In the pop-up window, look for the line that starts with **disk =**. The file names and locations appear on the same line. Scroll to the right to get the full name.
In this example, the virtual disk for vora1 is /OVS/Repositories/
0004fb0000030000a7e754fc1745a7c8/VirtualDisks/0004fb0000120000d55b46bbe6e73cfb.img



37. In the ssh session, navigate to the corresponding location on the view volume.

```
# cd /mnt/recover/VirtualDisks
# ls -l
total 976943104
-rw----- 1 root root 107374182400 Jan 15 21:40
    0004fb000012000011278ac725dba7aa.img
-rw----- 1 root root 107374182400 Jan 15 00:05
    0004fb000012000014612298c511d64d.img
-rw----- 1 root root 53687091200 Feb 6 2015
    0004fb00001200001aeba731eddafb87.img
-rw----- 1 root root 1073741824 Feb 13 2015
    0004fb0000120000252d3978cf2df87b.img
-rw----- 1 root root 53687091200 Feb 6 2015
    0004fb000012000027fbe59962cf14ca.img
-rw----- 1 root root 107374182400 Jan 4 16:00
    0004fb000012000032886294c9cbaf00.img
-rw-r--r-- 1 root root 44029706240 Sep 16 12:13
    0004fb0000120000371941ac56338d08.img
-rw-r--r-- 1 root root 107374182400 Dec 4 11:21
    0004fb00001200003c6a28102e3636f7.img
-rw-r--r-- 1 root root 107374182400 Dec 15 03:19
    0004fb00001200006c6757539e8ac6f6.img
-rw-r--r-- 1 root root 6720053760 Jan 6 09:48
    0004fb00001200006eae50e79962c798.img
```

```

-rw----- 1 root root 53687091200 Dec 15 10:19
0004fb00001200007a65875369599209.img
-rw----- 1 root root 107374182400 Dec 10 11:07
0004fb00001200007c8afa4929c34d09.img
-rw-r--r-- 1 root root 107374182400 Feb 6 2015
0004fb00001200008161448e7efab83c.img
-rw-r--r-- 1 root root 107374182400 Feb 6 2015
0004fb0000120000bf8d44e81ebfa8aa.img
-rw----- 1 root root 53687091200 Jan 7 11:15
0004fb0000120000ce0996781292f4d1.img
-rw----- 1 root root 107374182400 Jan 5 23:53
0004fb0000120000d55b46bbe6e73cfb.img
-rw-r--r-- 1 root root 107374182400 Dec 15 03:37
0004fb0000120000e28a56b2290c73b2.img
-rw-r--r-- 1 root root 107374182400 Dec 4 11:13
0004fb0000120000e3f641bff54780a3.img
-rw-r--r-- 1 root root 107374182400 Jan 6 16:22
0004fb0000120000f2af4760c4a8b8f1.img

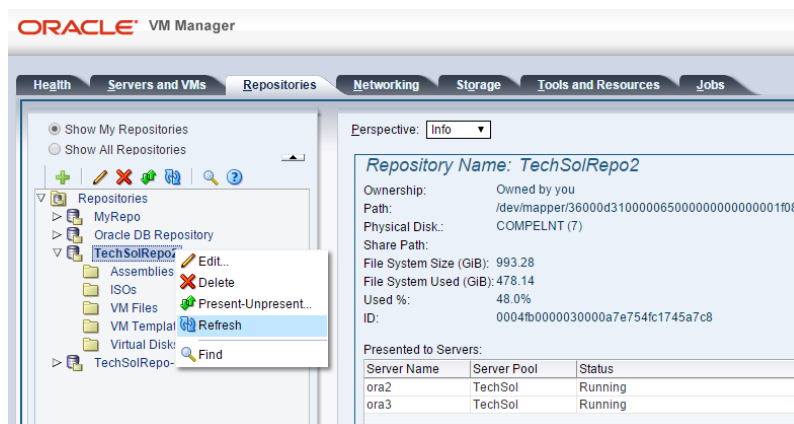
```

A point-in-time copy of the same virtual disk file is found on the view volume.


38. Restore the virtual disk back to the original repository. Make sure the virtual machine is shut down. Use the `cp` command to copy the file. This operation will overwrite the existing virtual disk so verify the information before running the command.

```
# cp -p 0004fb0000120000d55b46bbe6e73cfb.img \
/OVS/Repositories/0004fb0000030000a7e754fc1745a7c8/VirtualDisks
```

39. If the virtual disk was previously removed in Oracle VM Manager, rescan the repository to rediscover the restored virtual disk.



40. The restored virtual disk now shows up in the **Virtual Disks** folder. The restored disk might not have the original name so use the **File** column to confirm the existence of the disk. To rename the disk if desired, right-click the restored disk and select **Edit**.

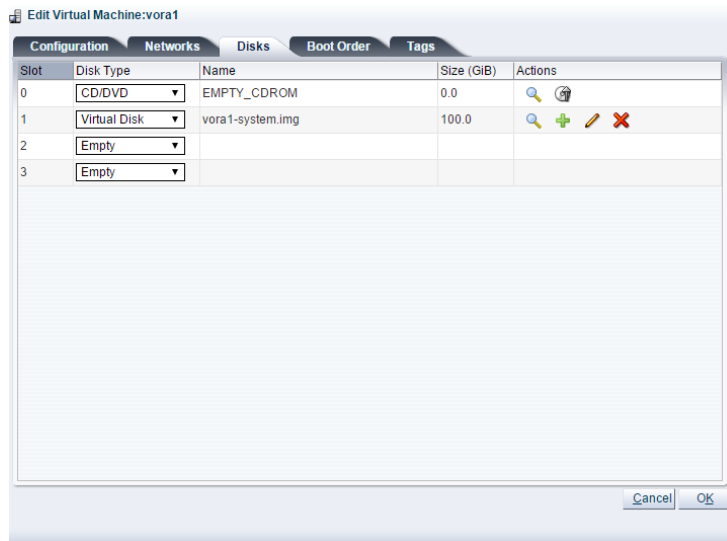


4.3.5 Updating the virtual machine that uses the virtual disk

If the virtual disk was removed from the virtual machine previously, reattach the virtual disk after it was restored to the virtual machine.

1. Under the **Servers and VMs** tab, select the server pool and change the **Perspective** to **Virtual Machines**.
2. Right-click the **Virtual Machine** in the right pane and select **Edit**.

41. Go to the **Disk** tab. Pick a new slot and choose the **Virtual Disk** disk type. Click the **Search** icon to add an existing virtual disk to the virtual machine.



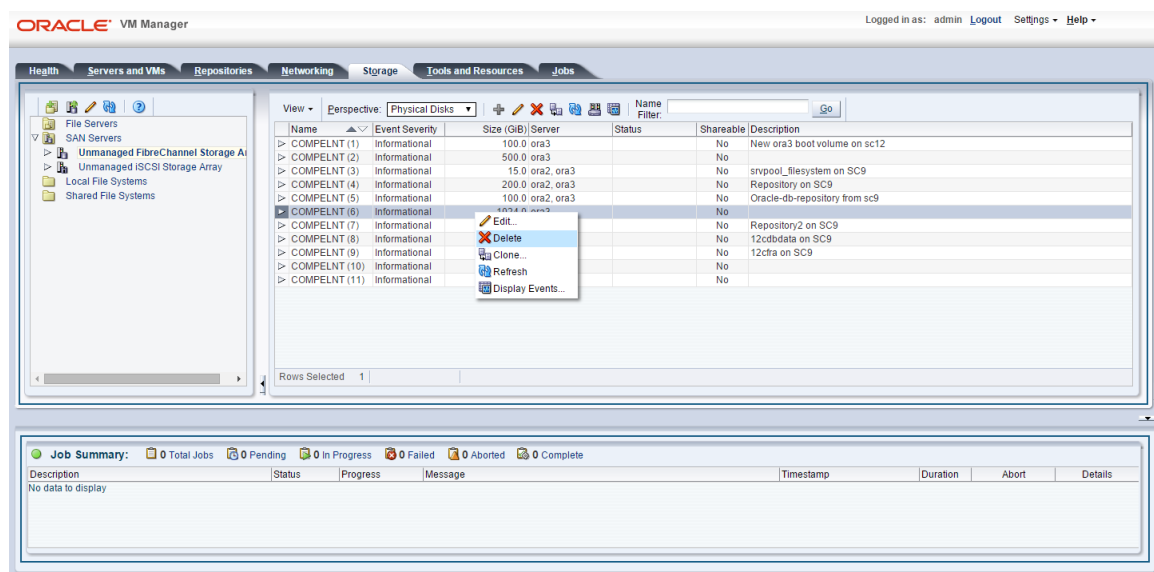
4.3.6 Cleaning up after restore

When the view volume is no longer needed, it should be removed from the Oracle VM and the SC array.

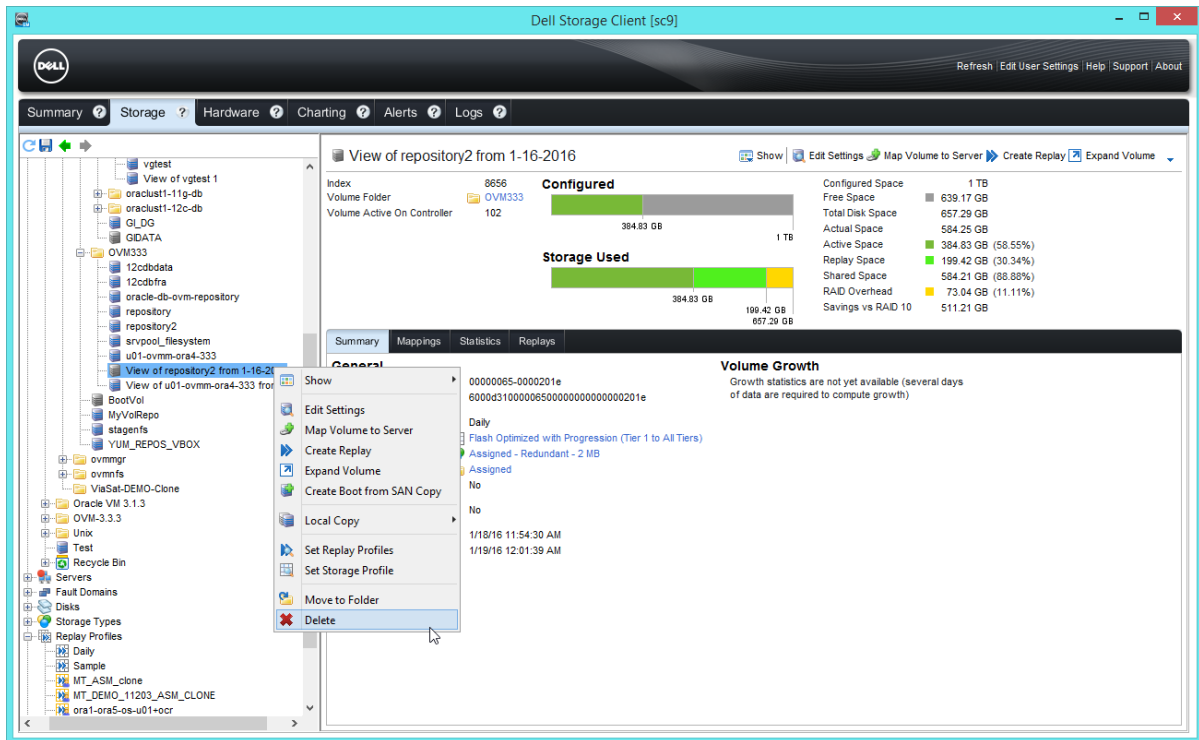
1. Unmount the file system.

```
# umount /mnt/recover
```

42. To remove the volume in the Oracle VM Manager web console, under the **Storage** tab, navigate to the SAN Servers, right-click the volume in the right pane, and select **Delete**. This removes the device from the Oracle VM Server and VM Manager.



43. Unmap the volume in Dell Enterprise Manager.
 - a. In Dell Enterprise Manager, under the **Storage** tab, right-click the volume and select **Remove Mappings**.
 - b. Right-click the volume again and select **Delete** to release the space back to the storage pool.



4.4 Restoring and recovering Oracle VM Manager

Oracle started shipping MySQL database in Oracle VM Manager 3.2 as the bundled database for its management repository. It is recommended to use the default bundled database for simple installation. Another advantage of using the bundled MySQL database is that it comes with a fully automated database backup and restore tool. The default backup file location is in `/u01/app/oracle/mysql/dbbackup`.

Oracle WebLogic Server also comes with automatic backup tool. The default backup files are located in `/u01/app/oracle/ovm-manager-3/domains/ovm_domain/servers/AdminServer/data/ldap/backup`.

Using Replay adds an extra layer of protection and provides fast recovery for everything residing on the SC volumes such as the Oracle VM Manager configuration file, application binaries, and backup files created by the Oracle backup tool.

As discussed in section 3.5, install Oracle VM Manager in a separate file system, normally in `/u01`, with one or more SC volumes so that they can be snapped and restored independently from the operating system. If multiple volumes are used, make sure they are configured in a **Consistent Replay** profile.

The following example replaces the file system /u01 from a Replay on the Oracle VM Manager server.

1. Stop the Oracle VM applications and database.

```
# service ovmm stop
# service ovmmcli stop
# service ovmm_mysql stop
```

44. Unmount the file system and disable the volume group.

```
# umount /u01
# vgchange -a n <vgname>
```

45. In Dell Enterprise Manager, remove the mapping of the volume(s).
46. Create a view volume for each volume used by the file system.
47. Map the volume(s) to the Oracle VM Manager server.
48. Scan and discover the new view volume(s) on the Oracle VM Manager server.

```
# /usr/bin/rescan-scsi-bus.sh
```

49. Reactive the volume group and mount the file system.

```
# vgchange -a y <vgname>
# mount /u01
```

50. Start the Oracle VM applications and database.

```
# service ovmm_mysql start
# service ovmm start
# service ovmmcli start
```

4.4.1 Mounting the Replay alongside the original file system

Instead of replacing the entire installation completely, restore just the files and components required.

1. Map the view volume(s) to the Oracle VM Manager server.
51. Scan and discover the new view volume(s).

52. Use the `vgimportclone` command to properly change the UUIDs of the cloned volume group and physical volume. Without doing so, the cloned volume group and file system cannot be mounted on the same system.

Use the `multipath` command to show all the multipath devices and their WWNs. Then compare the WWNs to the view volume device id. The WWN contains the view volume's device id. In this example, the original volume group is `vgoracle`. The cloned volume group is `vgoracle-snap`. The view volume device is `/dev/mapper/mpathy`.

```
# vgimportclone --basevgname vgoracle-replay /dev/mapper/mpathy
```

```
Physical volume "/tmp/snap.BETTr869/vgimport0" changed
1 physical volume changed / 0 physical volumes not changed
WARNING: Activation disabled. No device-mapper interaction will be
attempted.
Volume group "vgoracle" successfully changed
Volume group "vgoracle" successfully renamed to "vgoracle-replay"
Reading all physical volumes. This may take a while...
Found volume group "vg00" using metadata type lvm2
Found volume group "vgoracle-replay" using metadata type lvm2
Found volume group "vgoracle" using metadata type lvm2
Found volume group "vgnfs" using metadata type lvm2
Found volume group "vgyumvbox" using metadata type lvm2
```

53. Activate the cloned volume group and mount the file system under `/u01-snap`.

```
# vgchange -a y vgoracle-snap
# mkdir /u01-snap
# mount /dev/vgoracle-snap/lv_u01 /u01-snap

# df -h
Filesystem              Size  Used Avail Use% Mounted on
/dev/mapper/vg00-lv_root    50G   6.0G   41G   13% /
tmpfs                    32G     0   32G    0% /dev/shm
/dev/sdk1                 477M  103M  349M   23% /boot
/dev/mapper/vgoracle-lv_u01  19G   14G   4.8G   75% /u01
/dev/mapper/vgyumvbox-lvyumrepo  99G   79G   20G   80% /var/www/html/yum
/dev/mapper/vgyumvbox-lvvbox 296G  157G  139G   54% /vbox
/dev/mapper/vgnfs-lvnfs    591G  497G   94G   85% /stage
/dev/mapper/vgoracle--replay-lv_u01  19G   14G   4.9G   74% /u01-snap
```

Now both file systems are available. Copy the files from the restored file system if needed.

1 Terminology

Assembly: An assembly contains multiple virtual machine configurations including their virtual disks and connectivity settings between them. It is used to simplify the creation of 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.

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

OCFS2: Oracle VM uses OCFS2 (Oracle Clustered File System 2) to provide clustered and shared storage to Oracle VM Servers and virtual machines.

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

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

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

Replay: An SC Series array feature that provides a storage point-in-time copy of data, a Replay is used interchangeably with snapshot.

Replay profile: This defines the frequency and schedule for when Replays can run.

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

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

Storage repository: This 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 templates: These consist of a fully installed and configured virtual machine with all the required applications and software stacks pre-installed. Templates are reusable for creating new virtual machines with the exact image and configuration.

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

View volume: This 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: These are Fibre Channel zones used to segment the fabric to restrict access.

1 Technical support and resources

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.

1.1 Related documentation

The following documentation can be downloaded from the Knowledge Center on the SC Series [Portal](#) (login required):

- Dell Compellent Storage Center System Manager Version 6.x Administrator's Guide
- Dell Compellent Enterprise Manager 2015 R2 Administrator's Guide

Also see the following additional resources:

- [*Dell Storage Center with Red hat Enterprise Linux \(RHEL\) 6x Best Practices*](#)
- [*Oracle Virtualization Product Page*](#)
- [*Oracle VM Server for x86 Product Page*](#)
- [*Oracle Virtualization Product Documentation Libraries*](#)
- [*Oracle Database Documentation Library*](#)