

# Oracle VM 3 with Oracle ZFS Storage Appliance

Configuration Best Practices for Oracle ZFS Storage Appliance  
Using NFS v4 Protocol

ORACLE WHITE PAPER | DECEMBER 2015





## Table of Contents

Introduction	2
Configuring Oracle ZFS Storage Appliance for Oracle VM 3	3
Controller Settings and Software Releases	3
Storage Layout Recommendation for Oracle VM 3	3
IP Network Settings	5
Configuring Projects and Network File Systems for Oracle VM 3	8
Configuring Storage in the Oracle VM 3 Environment	14
Discovering Servers	14
Configuring Networking for NFS v4 Protocol	16
Configuring the Bonding Interface Property for the Servers	16
Configuring Storage Network Settings in Oracle VM Manager	17
Discovering the File Server to Establish Storage Communication	20
Configuring the Oracle VM 3 Cluster Server Pool with NFS Protocol	23
Configuring Oracle VM 3 Repositories	25
Maximizing Performance with Virtual Machine Setup	28
Virtual Disk Tuning for Oracle ZFS Storage Appliance	28
Overview of Tested System Components	30
Conclusion	31
References	32



## Introduction

This white paper provides best practices and recommendations for configuring Oracle VM 3 with Oracle ZFS Storage Appliance to reach optimal I/O performance and throughput. The outlined best practices and recommendations highlight configuration and tuning options for Oracle ZFS Storage Appliance NFS v4 protocol for Oracle VM 3.

The paper also includes recommendations for correct design of IP network infrastructure for Oracle VM 3, storage and virtual machine layout, as well as key settings to ensure high availability for Linux and Windows virtual machines running on top of Oracle VM 3 with Oracle ZFS Storage Appliance.

Highlighted in this paper are:

- Best practices and recommendations for deploying Oracle VM 3 with Oracle ZFS Storage Appliance using NFS v4 protocol
- Step-by-step configuration for Oracle VM 3 networking and server pools with Oracle ZFS Storage Appliance
- IP network design and tuning recommendation for NFS v4 protocol
- Oracle VM clustered server pool recommendations for high availability with Oracle ZFS Storage Appliance
- Oracle VM repositories and virtual machine data layout with Oracle ZFS Storage Appliance

NOTE: References to Sun ZFS Storage Appliance, Sun ZFS Storage 7000, and ZFS Storage Appliance all refer to the same family of Oracle ZFS Storage Appliances.

## Configuring Oracle ZFS Storage Appliance for Oracle VM 3

This section provides best practices and recommendations for Oracle VM 3 with Oracle ZFS Storage Appliance. The following configurations are recommended in order to optimize performance, throughput, and mainly to ensure high availability for Oracle VM environments using Network File System (NFS) v4 protocol with Oracle ZFS Storage Appliance.

### Controller Settings and Software Releases

For Oracle ZFS Storage Appliance controllers and software releases, the following minimum configuration is recommended for deployments in Oracle VM 3 environments:

- Oracle ZFS Storage Appliance ZS3-2, with at least 2x2 10GHz Intel® Xeon® CPU E5-2658 0 @ 2.10GHz Advanced Encryption Standard – New Instructions (AES-NI)
- Active/active cluster configuration mode
- Oracle ZFS Storage Appliance Software OS 8.4.2 or greater
- At least 512 GB of direct random access memory (DRAM) (L1 cache) per Oracle ZFS Storage Appliance head
- At least two 512 GB solid state drives (SSD) for ReadZilla per Oracle ZFS Storage Appliance head
- At least two dual 10 gigabit Ethernet (GbE) network interface cards (NICs) per Oracle ZFS Storage Appliance head, in order to be used by IP network multipathing (IPMP) configuration or bundled into a single 802.3ad Link Aggregation Control Protocol (LACP) channel

Figure 1 illustrates recommended configurations for Oracle ZFS Storage Appliance heads 1 and 2 for Oracle VM 3 environments with NFS v4 and iSCSI protocol.

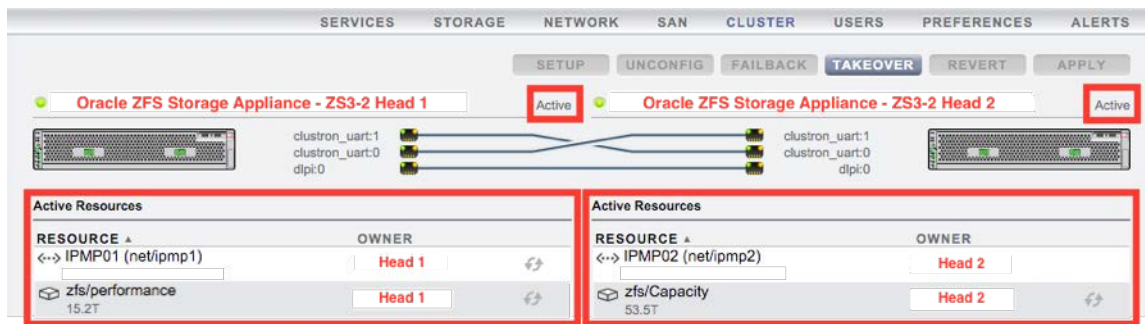


Figure 1. Oracle ZFS Storage Appliance – controller settings and software release recommended for Oracle VM 3 environments

### Storage Layout Recommendation for Oracle VM 3

Virtualized environments produce a high random I/O pattern, and will always require a storage device that is capable of providing low latency, fast response time, and also high availability. So to properly accommodate the high workloads demanded by critical applications running inside of the virtual machines, the recommended storage data profile for Oracle VM 3 deployments with Oracle ZFS Storage Appliance is **mirrored data profile**. Oracle ZFS Storage Appliance mirrored data profile will ensure and provide faster, reliable storage for critical applications running inside of the Oracle VM 3 virtual machines.

The minimum recommended storage layouts for Oracle VM 3 are:

- For Oracle VM 3 environments with Oracle ZFS Storage Appliance which demand more performance rather than capacity, a **mirrored data profile** disk pool of at least 44 x 900 GB (10000 RPM – Oracle ZFS Storage Appliance performance disks) with at least two 73 GB SSD devices for LogZilla (configured as striped log profile), and at least one 512 GB for L2 cache (L2ARC). The following figure illustrates this storage layout.



Figure 2. Oracle ZFS Storage Appliance – recommended storage layout for Oracle VM 3 performance environments

- For Oracle VM 3 environments with Oracle ZFS Storage Appliance which demand more capacity rather than performance, a **mirrored data profile** disk pool of at least 44 x 4 TB (7200 RPM – Oracle ZFS Storage Appliance capacity disks) with at least two 73 GB SSD devices for LogZilla (configured as striped log profile), and at least one 512 GB for L2 cache (L2ARC), as seen in figure 3, is recommended.




Figure 3. Oracle ZFS Storage Appliance – recommended storage layout for Oracle VM 3 capacity environments

- For Oracle VM 3 environments with Oracle ZFS Storage Appliance which demand both performance and capacity, a **mirrored data profile** disk pool of at least 20 x 900 GB (10000 RPM – Oracle ZFS Storage Appliance performance disks) with at least two 73 GB SSD devices for LogZilla (configured as

striped log profile), and at least one 512 GB for L2 cache (L2ARC), and also a **mirrored data profile** disk pool of at least 44 x 4 TB (7200 RPM – Oracle ZFS Storage Appliance capacity disks) with at least two 73 GB SSD devices for LogZilla (configured as striped log profile), and at least one 512 GB for L2 cache (L2ARC) are recommended. This storage layout will ensure capacity, performance, and also high availability for Oracle VM 3 environments with Oracle ZFS Storage Appliances. See figure 4.

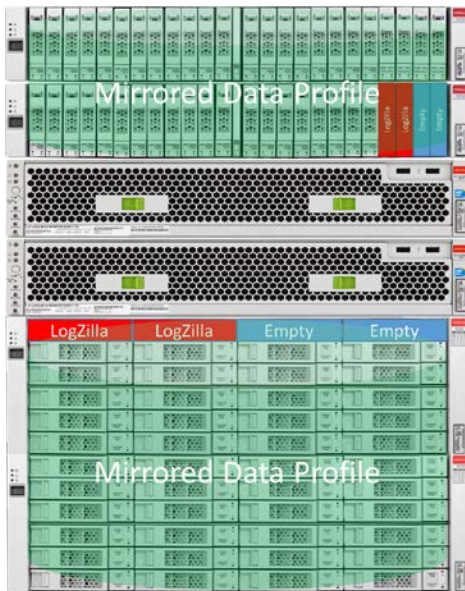



Figure 4. Disk pool configuration for Oracle ZFS Storage Appliance with Oracle VM 3 and maximized capacity, performance, and high availability environments

## IP Network Settings

For network configuration, both IPMP and LACP IEEE 802.3ad bonding mode are recommended for deployment of Oracle VM 3 using NFSv4 protocol with Oracle ZFS Storage Appliance. IP network multipathing protocol (IPMP) provides physical interface failure detection, increased reliability, and mainly network performance for systems with multiple interfaces. LACP IEEE 802.3ad provides a method that combines the capacity of multiple full-duplex Ethernet links into a single logical link, which increases bandwidth by aggregating existing datalinks, load balance of inbound and outbound network traffic, and also automatic failover/failback traffic from a failed link over to working links in the aggregation. These two technologies complement each other, and can be deployed together to provide benefits for network performance, bandwidth, and high availability for Oracle VM 3 environments with Oracle ZFS Storage Appliance.

For the LACP configuration, ensure that the Oracle ZFS Storage Appliance 10GbE network interfaces, IP switches, and also Oracle VM 3 servers are correctly configured with the following LACP policies:

- LACP policy L3 – Oracle ZFS Storage Appliance and IP network switches
- LACP active mode – Oracle ZFS Storage Appliance and IP network switches
- LACP short timer interval – Oracle ZFS Storage Appliance and IP network switches
- Link Aggregation Policy – Oracle VM 3 servers network bonding configuration



For IPMP configuration, ensure that the network bonding configuration in the Oracle VM 3 servers is correctly set with Load Balanced policy.

For all 10GbE network interfaces used by the Oracle VM 3 environment, ensure that 9000 maximum transmission unit (MTU) jumbo frame is enabled on Oracle ZFS Storage Appliance 10GbE network interfaces, IP switches, and also on Oracle VM 3 servers.

To test for 9000 MTU jumbo frame, you can use `ping` or `tracepath` Linux commands in the Oracle VM servers, as the following example shows:

```
ping -M do -s 8972 <Oracle ZFS Storage Appliance 10GbE interface IP address>
```

```
PING x.x.x.x (x.x.x.x) 8972(9000) bytes of data.  
8980 bytes from x.x.x.x: icmp_seq=1 ttl=255 time=0.162 ms  
8980 bytes from x.x.x.x: icmp_seq=2 ttl=255 time=0.165 ms  
8980 bytes from x.x.x.x: icmp_seq=3 ttl=255 time=0.141 ms  
8980 bytes from x.x.x.x: icmp_seq=4 ttl=255 time=0.164 ms
```

Example of `tracepath` Linux command:

```
tracepath <Oracle ZFS Storage Appliance 10GbE interface IP address>
```

```
1: x.x.x.x (x.x.x.x) 0.088ms pmtu 9000
```

```
Resume: pmtu 9000 hops 1 back 255
```

Figure 5 presents a sample of an IPMP network configuration recommended for Oracle ZFS Storage Appliance (head 1) within an Oracle VM 3 environment.



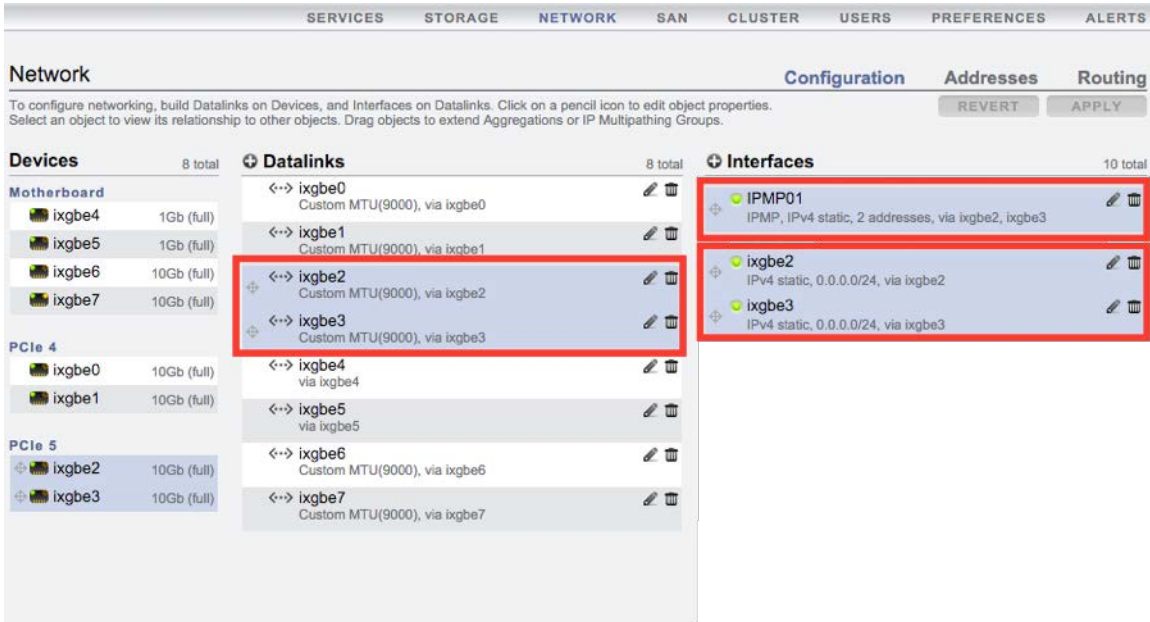


Figure 5. Example of IPMP network configuration on Oracle ZFS Storage Appliance cluster head 1 for Oracle VM 3

Figure 6 shows a sample recommended IPMP network configuration for Oracle ZFS Storage Appliance (head 2) within an Oracle VM 3 environment.

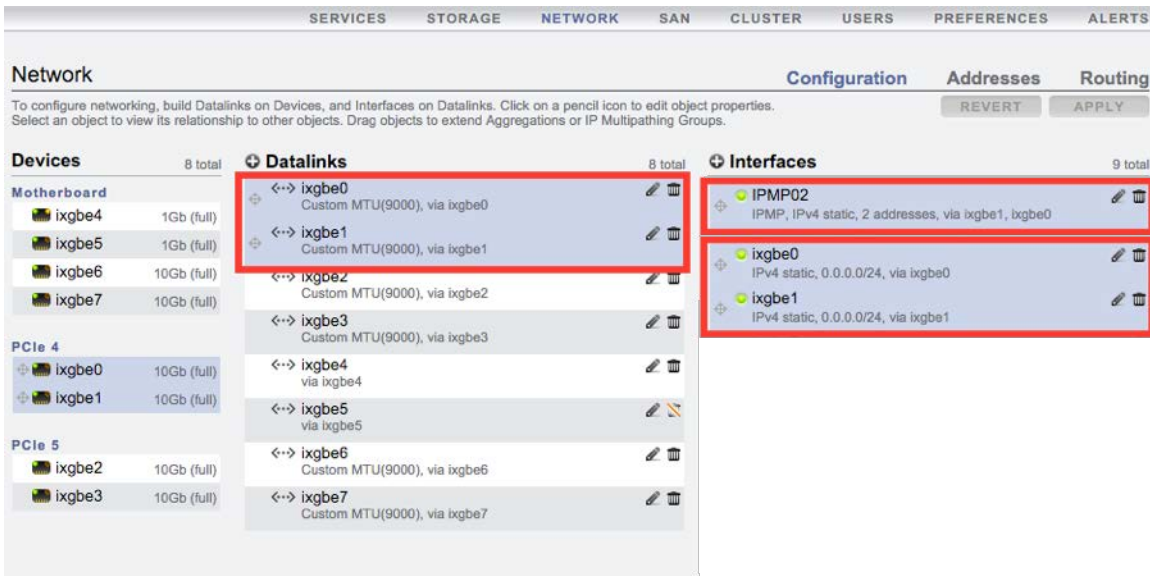


Figure 6. Example IPMP network configuration for Oracle ZFS Storage Appliance cluster head 2 and Oracle VM 3

Figure 7 shows a recommended network interface configuration on the Oracle ZFS Storage Appliance for IPMP. To have at least two 10GbE network interfaces as part of the IPMP configuration in active/active mode, or bundled into a single 802.3ad LACP channel, will require at least two dual 10GbE network interface cards per Oracle ZFS Storage Appliance head.



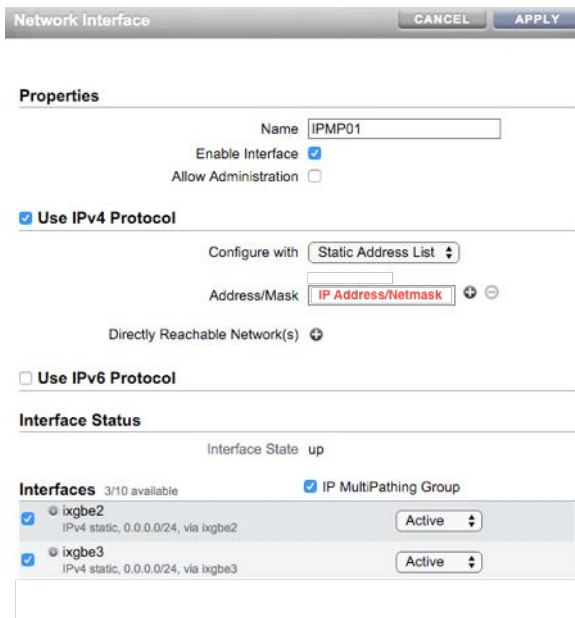


Figure 7. IPMP network configuration on Oracle ZFS Storage Appliance cluster head 2 for Oracle VM 3

## Configuring Projects and Network File Systems for Oracle VM 3

This section presents the best practices and recommendations for setting up Oracle ZFS Storage Appliance projects, NFS shares, and Oracle VM repositories.

As a best practice for better performance and high availability of Oracle ZFS Storage Appliance NFS shares presented to Oracle VM environment, split the NFS shares across different cluster heads and different Oracle ZFS Storage Appliance storage pools. This model allows specific, tailored configuration for the NFS shares – such as different file sizes, compression, and caching – which is critical to meeting the requirements of the application running inside of the virtual machines.

The recommended storage layout of Oracle ZFS Storage Appliance for Oracle VM 3 repositories:

- Set up one NFS share per Oracle VM repository.
- On the Oracle ZFS Storage Appliance side, always match the database recordsizes to the requirements of the applications running inside of the virtual machines.
- Split the NFS shares across both heads of the Oracle ZFS Storage Appliance cluster.
- For critical applications which require low latency and high performance, use Oracle ZFS Storage Appliance performance disks – 10k RPM disks.
- For best performance, ensure that 9000 Maximum Transmission Unit (MTU) has been correctly enabled.

Figure 8 illustrates the layout and mapping for NFS shares in the Oracle ZFS Storage Appliance for Oracle VM 3 environments and repositories.

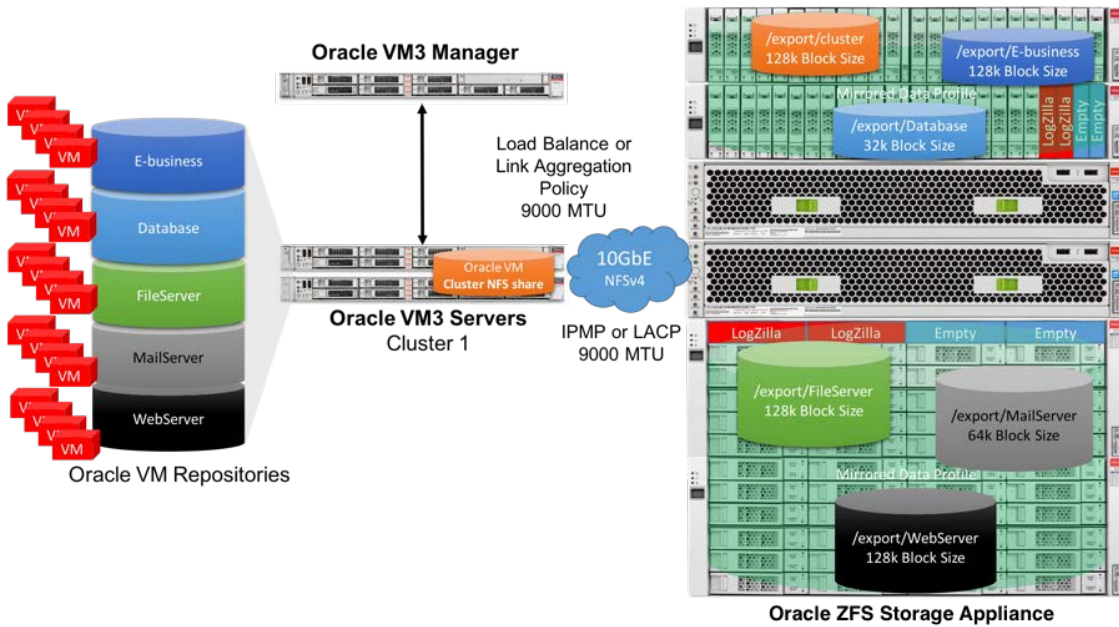


Figure 8. Oracle ZFS Storage Appliance NFS shares and Oracle VM repository layout

A new project for Oracle VM 3 should be created in the Oracle ZFS Storage Appliance BUI, as seen in the following figure. As seen in the highlighted text labels, click Shares, Projects, and then click the plus signal (+). Enter the name of the project.


NOTE: The example shown uses Oracle VM3 as the project name for the Oracle VM environment. All Oracle ZFS Storage Appliance NFS shares that will be presented later on to the Oracle VM environment, and also configured as Oracle VM repositories, will be configured under this OracleVM3 project.

NOTE: Encryption has not been used by the following example.



Figure 9. Configuring a new project in the Oracle ZFS Storage Appliance BUI

After successfully creating the OracleVM3 project, and assuming that an Oracle VM cluster server pool will be created, the first network file system to be created is the cluster file system. The cluster file system will be presented to the Oracle VM servers, and then configured as a server pool file system. A server pool file system is a storage area used by Oracle VM servers to hold the server pool, cluster data, and also for cluster heartbeating across the Oracle VM servers.



NOTE: Oracle recommends that the server pool file system be created with at least 12 GB for size. As a best practice for Oracle ZFS Storage Appliance configuration, use the following settings for a cluster file system of at least 15 GB of size:

- No data compression, or encryption enabled
- Database record size set to 128k
- NFS exceptions correctly configured with access granted only to the Oracle VM servers
- For security purposes, NFS share with 700 permission (only user has permissions)
- Services such as: SMB, HTTP, FTP, SFTP, TFTP disabled for file system. Only NFS service should be enabled as read/write share mode.
- File system properties such as: read only, update access time on read, non-blocking mandatory locking, data deduplication, and virus scan disabled
- Checksum property set to Fletcher4 (standard)
- Cache device usage set to all data and metadata
- Synchronous write bias set to latency
- Additional replication set to normal
- File system properties such as prevent destruction and restrict ownership change enabled
- File system quota and reservation enabled and set to 15 GB of size
- Default configuration accepted for the Access, Snapshot and Replication tabs

Figure 10 presents the initial configuration of an Oracle ZFS Storage Appliance network file system – including project name, file system name, permission, and standard configurations such as normalization, encryption, and mount point – which will later be used as a cluster file system for the Oracle VM 3 environment.

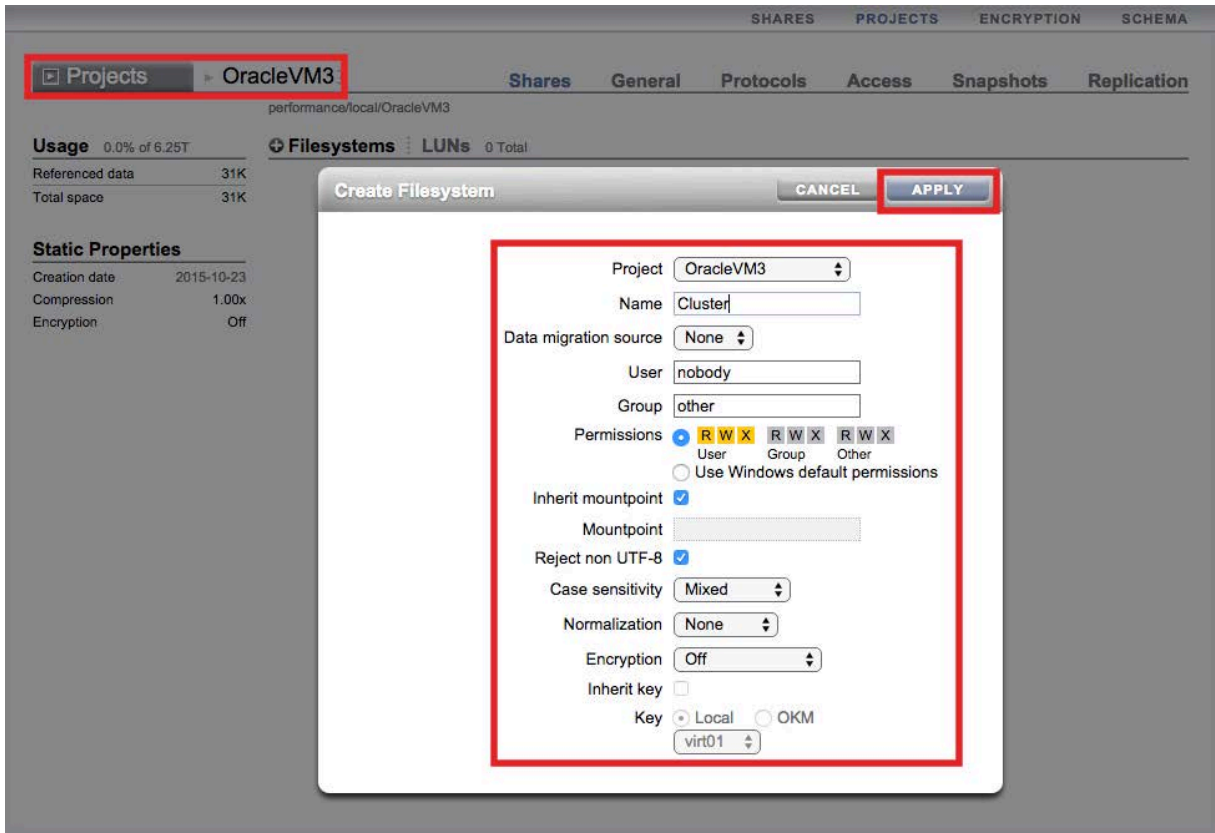


Figure 10. Project configuration shown in the Oracle ZFS Storage Appliance BUI

Figure 11 presents the general properties of the network file system. As previously recommended, the cluster file system should have quota enabled and reservation enabled, and the disk space should be set to at least 15 GB of size.

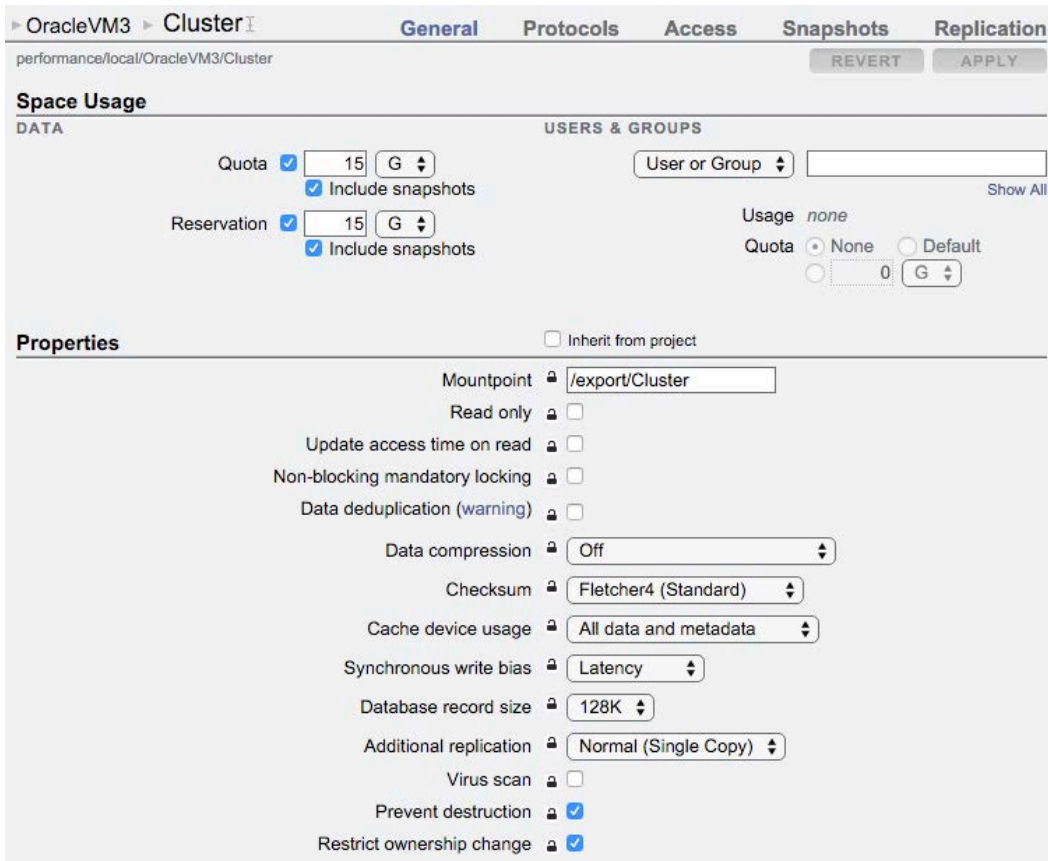


Figure 11. File system general properties shown in Oracle ZFS Storage Appliance BUI

Figure 12 presents the protocols and service status for the cluster file system. As a best practice, NFS service should be enabled with appropriated NFS exceptions, and share mode should be set to read/write.

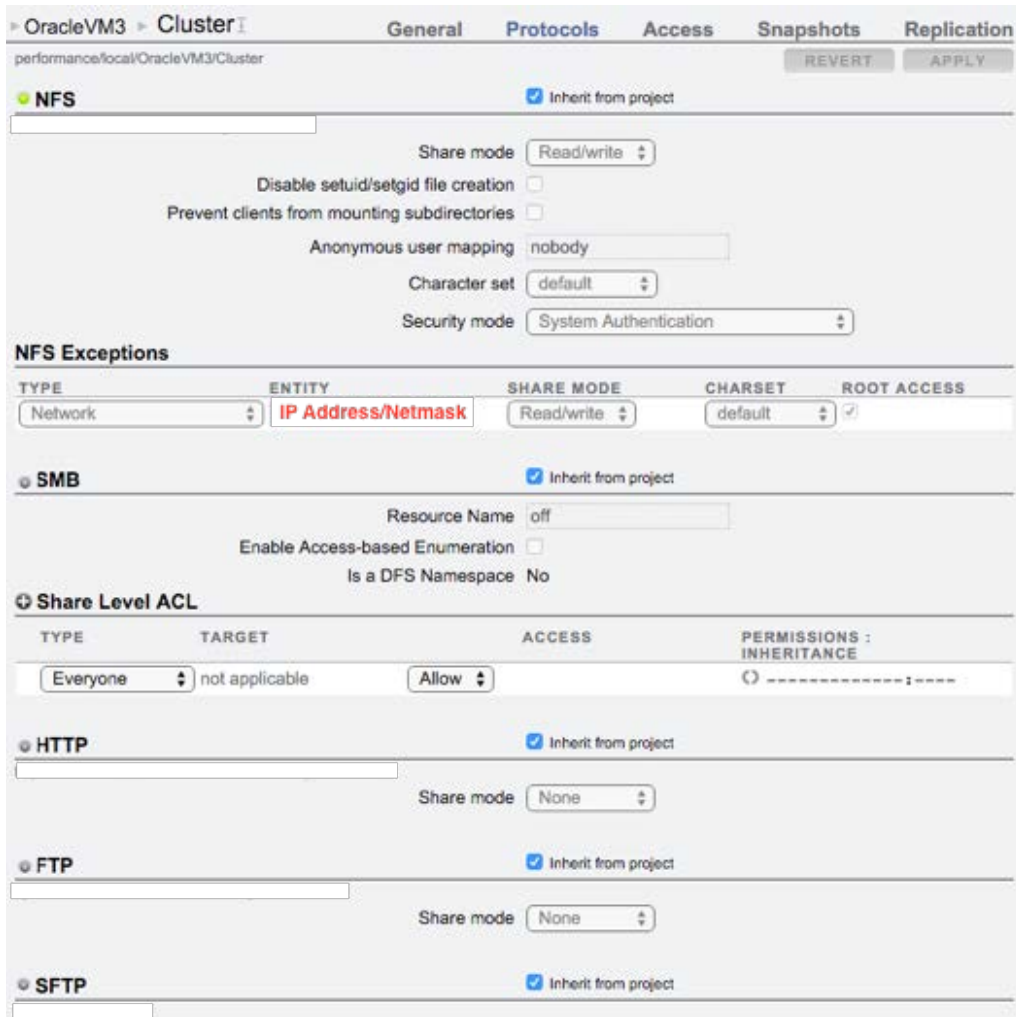


Figure 12. File system protocol status shown in the Oracle ZFS Storage Appliance BUI

As a best practice for Oracle ZFS Storage Appliance with Oracle VM 3, all file systems that will be presented as repositories to Oracle VM 3 servers should have the following configuration:

- Data compression enabled: LJZB
- Database record size set according to application requirements. Use these best practice guidelines:
  - For OLTP databases – 32k
  - Applications such as E-business suite, Customer Relationship Management (CRM), WebServers and FileServer – 128k
  - Mail server applications such as Microsoft Exchange Server – 64k
  - Repositories for Oracle LINUX and Windows Server virtual machines – 128k
- NFS exceptions correctly configured, with access granted only to the Oracle VM servers
- For security purposes, NFS share set with 700 permission

- Services such as SMB, HTTP, FTP, SFTP, and TFTP disabled for a file system presented to the Oracle VM environment. Only NFS service should be enabled as read/write share mode.
- Filesystem properties such as: read only, update access time on read, non-blocking mandatory locking, data deduplication and virus scan disabled
- Checksum property set to Fletcher4 (standard)
- Cache device usage set to all data and metadata
- Synchronous write bias set to latency
- Additional replication set to normal
- Filesystem properties such as prevent destruction and restrict ownership change enabled
- Access, Snapshot and Replication tabs left as the default configuration
- Filesystem quota and reservation should not be set for Oracle ZFS Storage Appliance file systems which will be configured as Oracle VM repositories and also presented to Oracle VM servers and virtual machines.

Figure 13 shows the Oracle ZFS Storage Appliance file systems configured with the settings previously listed. In the following section of this paper, these file systems will be presented to Oracle VM 3 servers and configured as repositories.

NAME ▲	SIZE	MOUNTPOINT	ENCRYPTED
Database	38K	/export/Database	
E-Business	38K	/export/e-business	
FileServer	9.54G	/export/FileServer	
MailServer	161G	/export/MailServer	
WebServer	26.6G	/export/WebServer	

Figure 13. Oracle ZFS Storage Appliance filesystems configured for Oracle VM shown in Oracle ZFS Storage Appliance BUI

## Configuring Storage in the Oracle VM 3 Environment

This session explains how to present Oracle ZFS Storage Appliance network file systems to the Oracle VM environment. Included are recommendations for configuring a storage network and server pool using Oracle ZFS Storage Appliance NFS, and for creating an Oracle VM clustered server pool on top of the Oracle ZFS Storage Appliance network file system.

### Discovering Servers

Assuming that the Oracle VM Manager and servers have been correctly installed and configured, the first step is to discover the Oracle VM servers. As highlighted on the Oracle VM Manager console shown in the following figure, click on the Servers and VMs tab, and then click on the Discover Server icon.



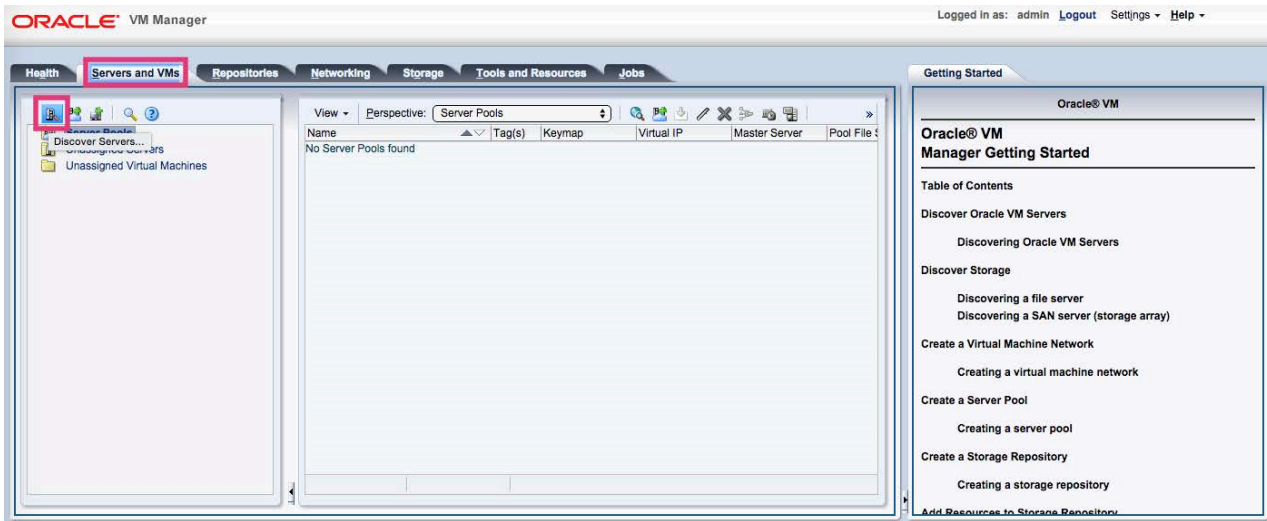


Figure 14. Selecting the Oracle VM discover server function in the Oracle VM Manager console

The next Oracle VM screen will prompt for Oracle VM Agent credentials such as password and IP addresses/DNS hostnames of the Oracle VM servers.

NOTE: The Oracle VM agent password used here is the same one prompted for during the Oracle VM servers installation.

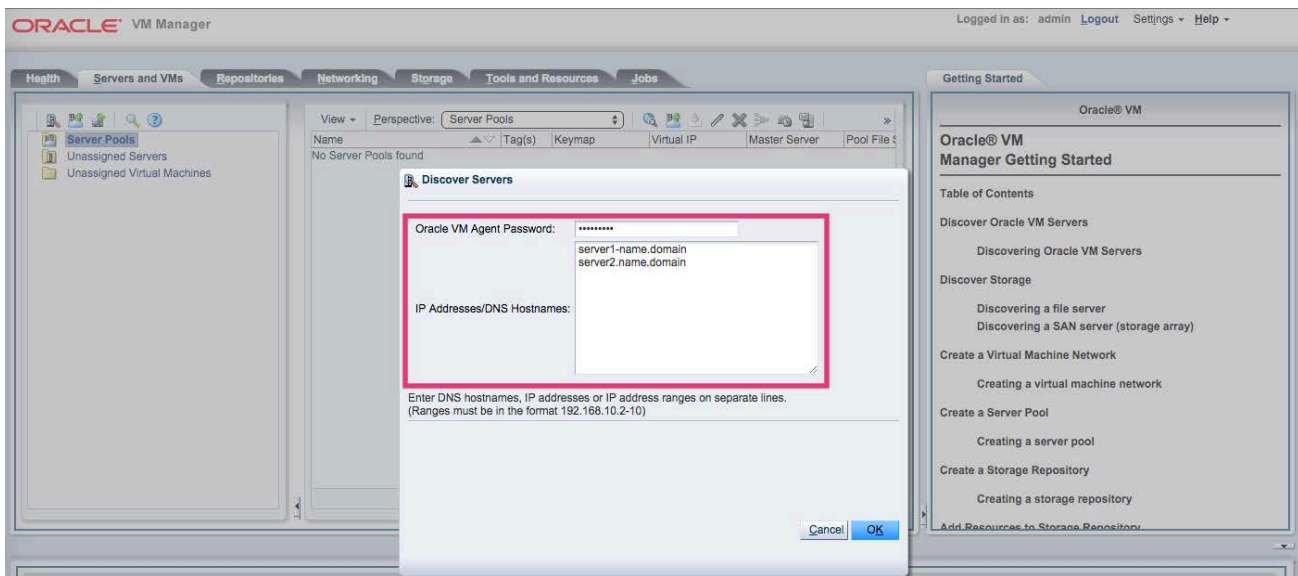


Figure 15. Oracle VM discover servers dialog window prompting for identifying agent and host information

After clicking OK, the Discover Servers operation should complete and display the Oracle VM servers in the Unassigned Servers folder, so they are available for network and server pool configuration.

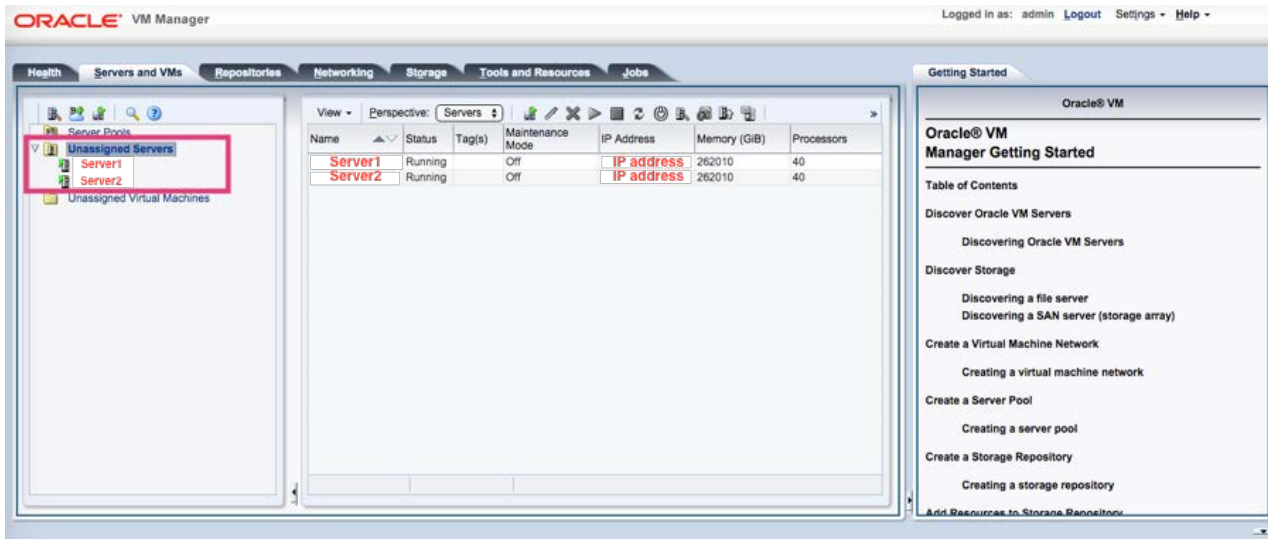


Figure 16. Oracle VM servers shown in the Oracle VM Manager console's Unassigned Servers folder before server pool configuration

## Configuring Networking for NFS v4 Protocol

Assuming that the Oracle VM servers have been successfully discovered and are available under the Unassigned Servers folder in Oracle VM Manager console, the next step is the configuration of the 10GbE network infrastructure, which will be utilized by Oracle VM servers to establish communication with Oracle ZFS Storage Appliance network file systems. This network is referred to as the storage network.

The recommended storage network setup for Oracle VM servers is:

- At least one dual 10GbE network interface card per Oracle VM server, in order to be used by Load Balance or bundled into a single 802.3ad LACP channel
- An isolated network, which can be accomplished utilizing VLANs, network segmentation, or dedicated IP switches for NFS traffic only
- Maximum transmission unit (MTU) jumbo frame (9000 bytes) correctly configured on Oracle VM servers, IP network switches, and Oracle ZFS Storage Appliance

However, before configuring the storage network, ensure that each Oracle VM server that will be part of the Oracle VM clustered server pool has the bonding interface property configured.

## Configuring the Bonding Interface Property for the Servers

Individually set the bonding interface in the Oracle VM Manager console for the servers by clicking on the Servers and VMs tab. Select a server listed in the Unassigned Servers folder, then click on the Perspective list box, where you can select the Bond Ports option. Select the (+) sign next to Create Bond Port to add a Bond network port for the Oracle VM server, and provide the following configuration information in the Create Bond Port dialog window illustrated in the following figure:

- IP address – IP address of the Oracle VM server Bond Port
- Mask – The subnet mask

- MTU – Maximum Transmission Unit, jumbo frame; must be set to 9000
- Description – A brief description for this Bond interface
- Bonding – If working with LACP IEEE 802.3ad, select Link Aggregation bonding mode. For IPMP, select Load Balance bonding mode.
- Selected Ports – Select the 10GbE network interfaces which will be used for the storage network.

Create the defined Bond Port by clicking the OK button.

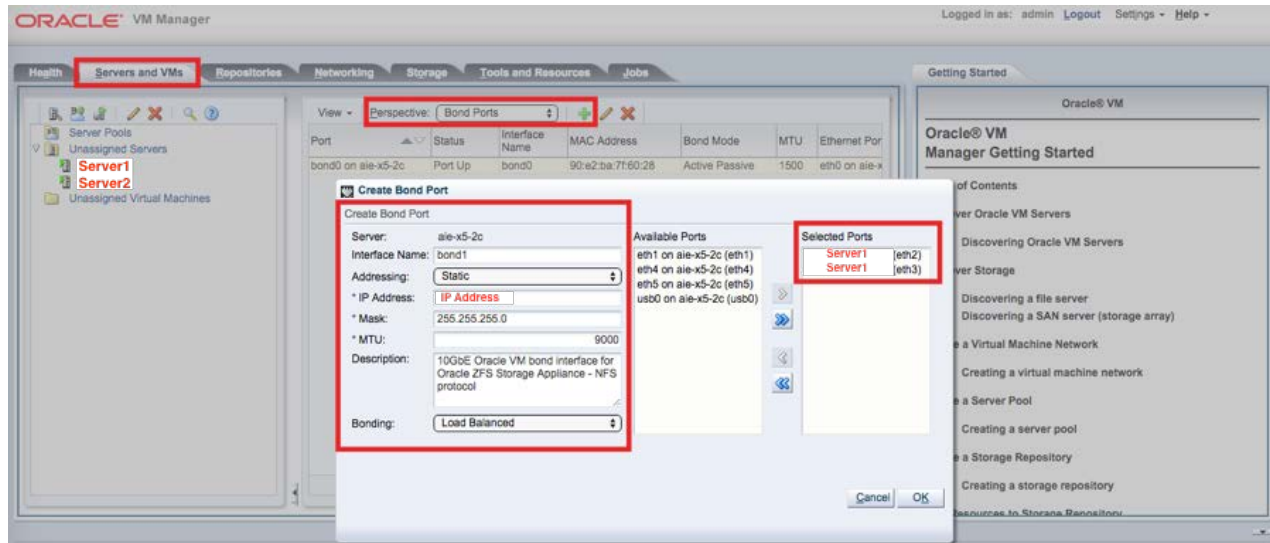


Figure 17. Bond Port configuration shown in Oracle VM Manager console

In this example, eth2 and eth3 are the 10GbE network interfaces that already are property configured and connected to an isolated 10GbE network or VLAN connected to the 10GbE IP switch ports. The switch ports, in turn, already have correct settings such as: 9000 MTU jumbo frame, LACP IEEE 802.3ad, or IPMP.

For best performance, all Oracle VM server members of the Oracle VM cluster server pool require at least one bond interface with the same bonding mode, which can be LACP IEEE 802.3ad or IPMP, and correctly configured with 9000 MTU jumbo frame.

NOTE: Active/backup bonding mode can be used in very few cases, which are out of the scope of this document. For best performance, resiliency, and also high availability with Oracle ZFS Storage Appliance, LACP IEEE 802.3ad or IPMP are the recommended bonding modes for the storage network with Oracle ZFS Storage Appliance.

Once the bond interfaces are correctly configured in all Oracle VM servers, you can proceed to storage network configuration. All storage communication between Oracle VM servers and Oracle ZFS Storage Appliance will be using this 10GbE storage network.

### Configuring Storage Network Settings in Oracle VM Manager

In the Oracle VM Manager console, click on the Networking tab, then click on the add (+) sign, or Create Network icon, under the Networks label. In the resulting pop-up window, choose the Create a Network with Ports/Bond Ports/VLAN Interfaces option, then click the Next button, as highlighted in the following figure.

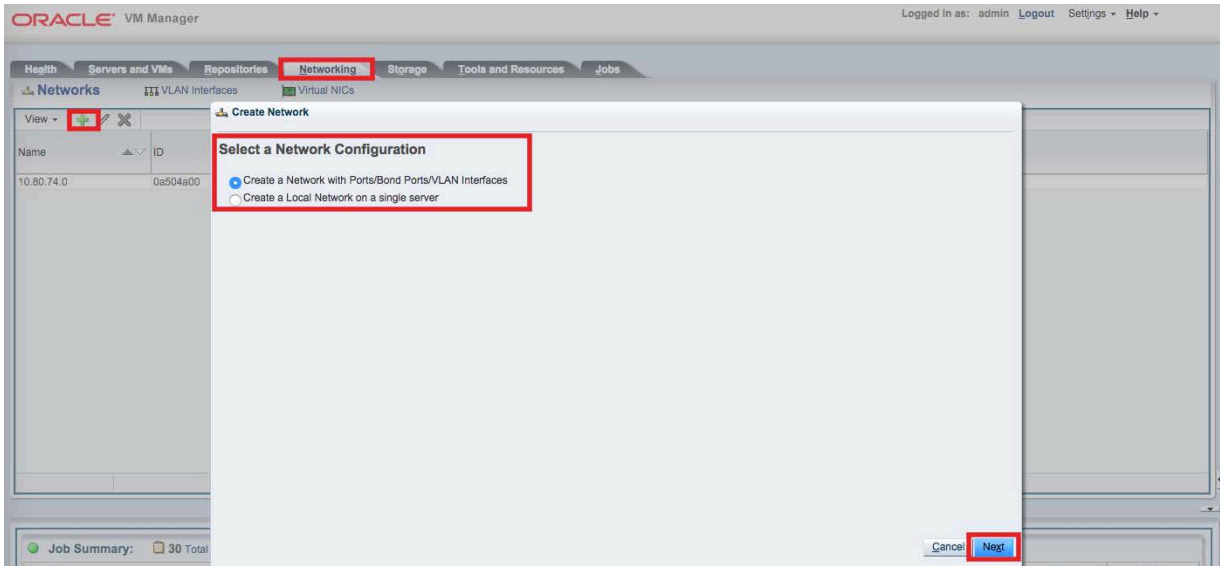


Figure 18. Defining the storage network in the Oracle VM Manager console

Following the screen prompts, provide the network name and a brief description for this network and select the Live Migrate and Storage checkmark options available under the Network Uses column. Click Next.

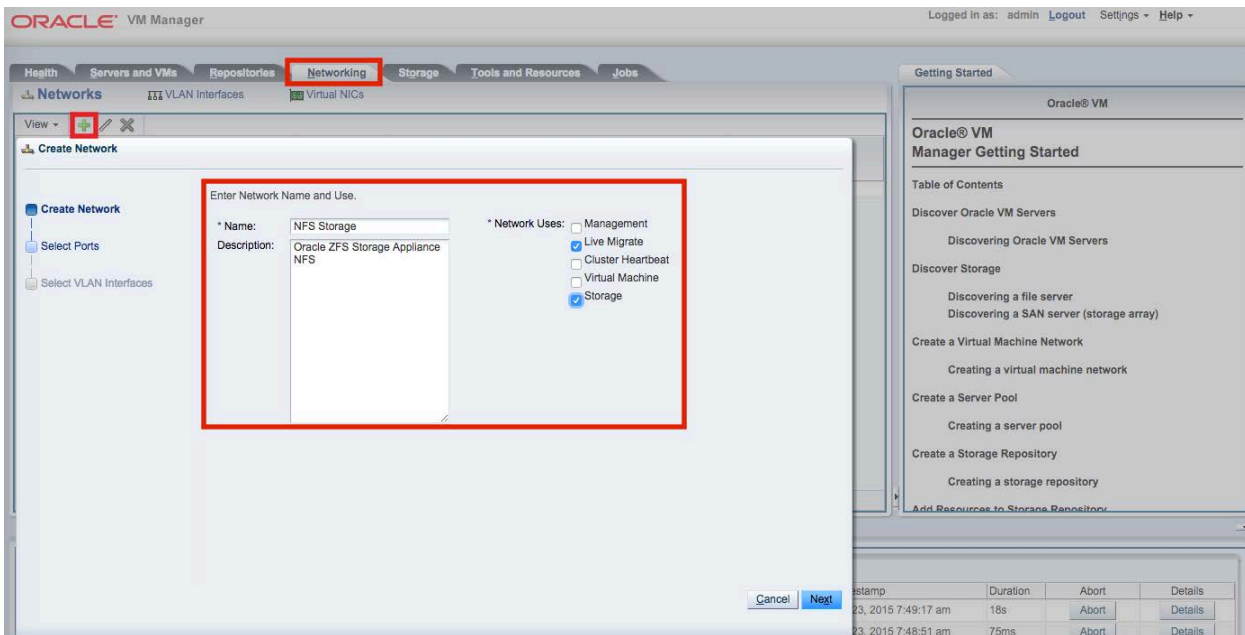


Figure 19. Naming and describing the storage network in the Oracle VM Manager console

The next screen will prompt for the ports to add to the network. Select the bond ports which were previously configured. In this example shown in the following figure, the bond1 port is selected. Click Next.

NOTE: The following screenshot reflects that only two Oracle VM servers have been used by the example solution.

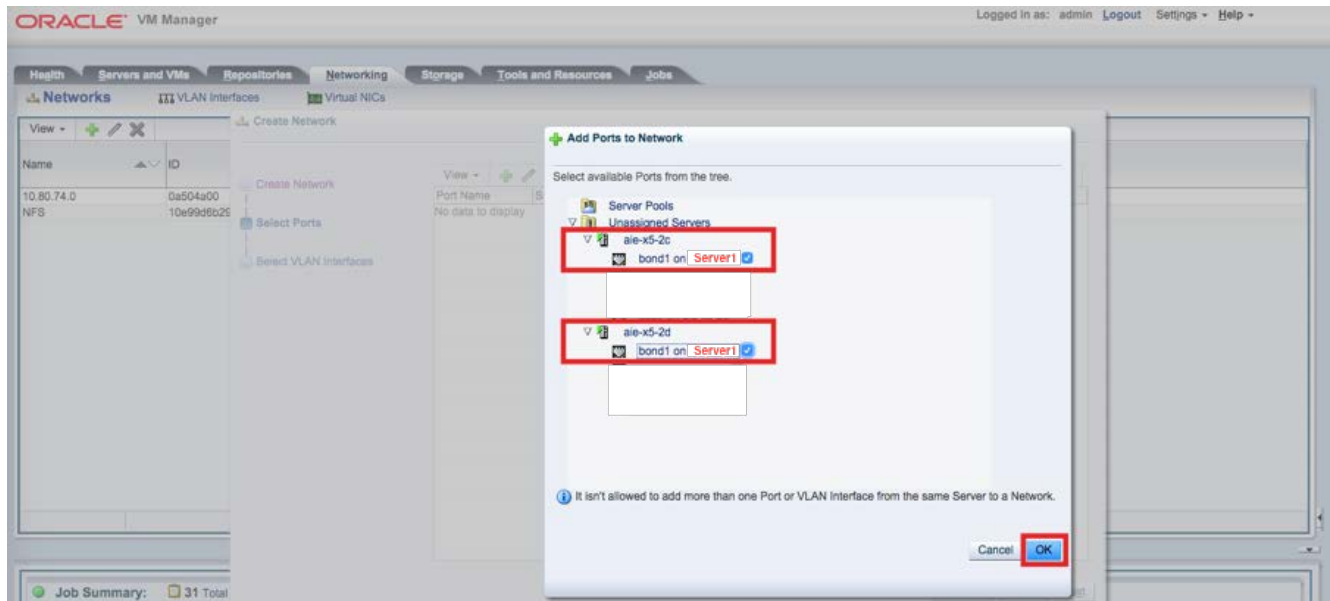


Figure 20. Adding ports for the storage network in the Oracle VM Manager console

The next screen will display all the bond ports. Ensure that they are configured for 9000 MTU and Bonding property set (shown as Yes), as seen in the following figure. Click Next.

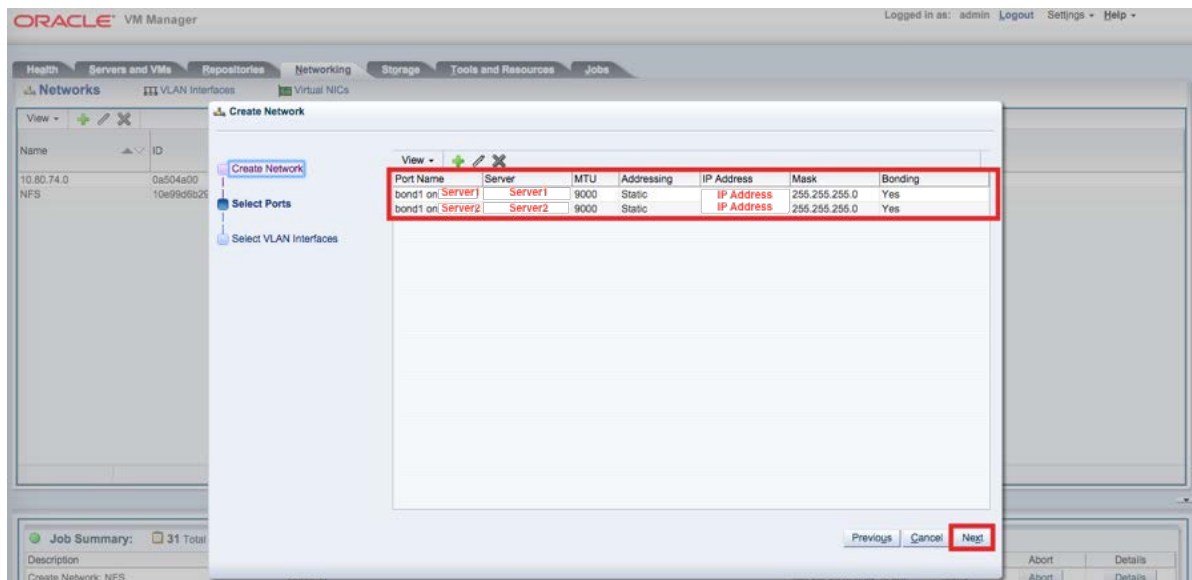


Figure 21. Checking the Bond Port configuration in the Oracle VM Manager console

As you can see from the figure and the Create Network dialog window, configuration includes a final step, 'Select VLAN Interfaces.' Assuming that only network isolation has been used for this document, VLAN interface configuration can be skipped. However, if the storage network is part of a different VLAN, the VLAN interfaces should be correctly configured and presented to the Oracle VM Servers.

Once you have gone through the Create Network dialog windows and committed your choices by selecting OK, the new storage network will be listed under the Networking tab in the Oracle VM Manager console, as shown in the following figure.

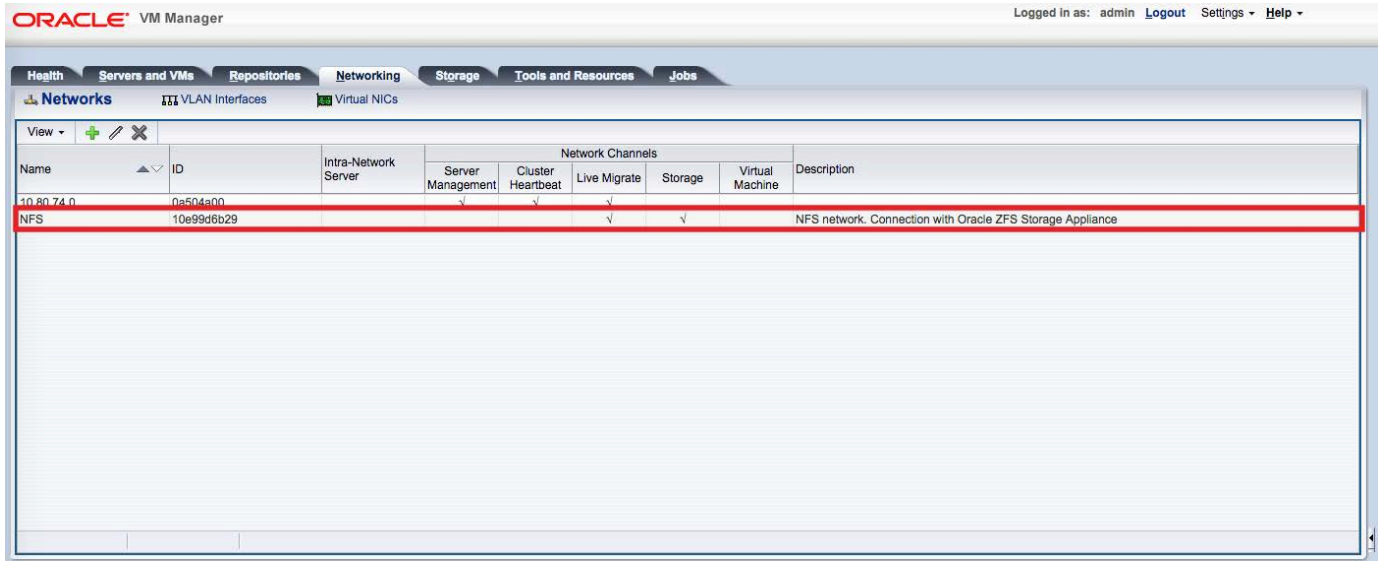


Figure 22. Successfully configured Bond Port displayed in the Oracle VM Manager console

### Discovering the File Server to Establish Storage Communication

After the storage network has been successfully configured, storage communication between Oracle VM servers and the Oracle ZFS Storage Appliance must be established in order to present Oracle ZFS Storage Appliance network file systems to Oracle VM servers. Engaging the Discover File Server option in the Oracle VM Manager console accomplishes this task. In the Oracle VM Manager console, click the Storage tab, then click Discover File Server, as highlighted in the following figure.

The Discover File Server operation will engage a plug-in to help establish the communication link. Assuming that NFS protocols have been used, the Oracle VM Generic Plug-in will be used.




Figure 23. Discovering the file server in the Oracle VM Manager console

The first screen, seen in the following figure, will prompt for:



- The storage plug-in to use to establish communication with the file server; in this case, Oracle Generic Network File System (the Oracle VM Generic Plug-in).
- The file server name (the Oracle ZFS Storage Appliance head 1 name)
- The access host IP address (the Oracle ZFS Storage Appliance 10GbE Link – LACP IEEE 802.3ad or IPMP)
- A brief description of the file server (the Oracle ZFS Storage Appliance head 1).

Select the Uniform Exports checkbox option, as shown. Select Next to proceed to the next input screen.

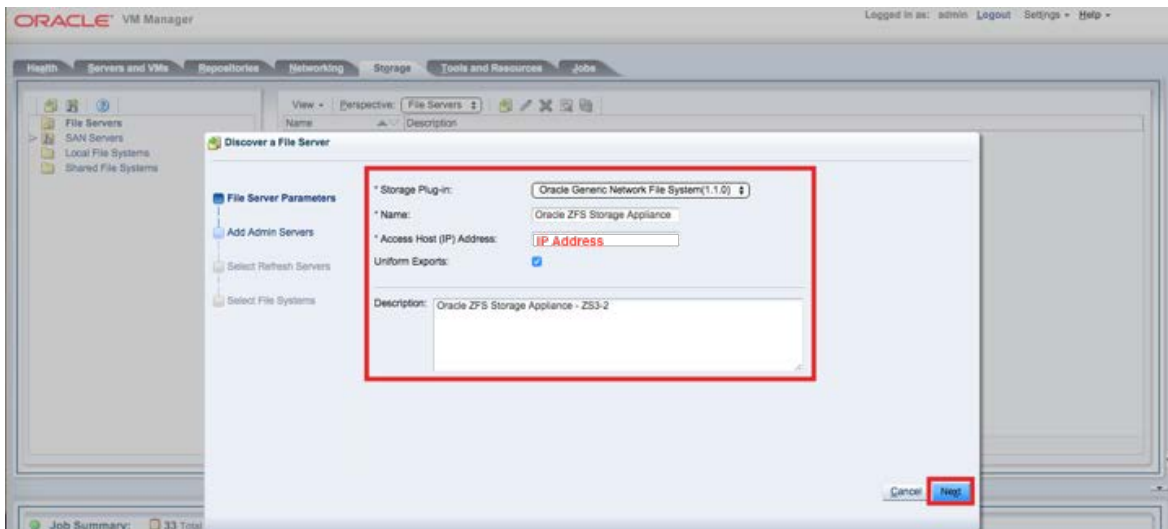


Figure 24. Specifying the storage plug-in and file server in the Oracle VM Manager console

In the next dialog window, select the admin servers that will have administrative access to this file server. Click Next.

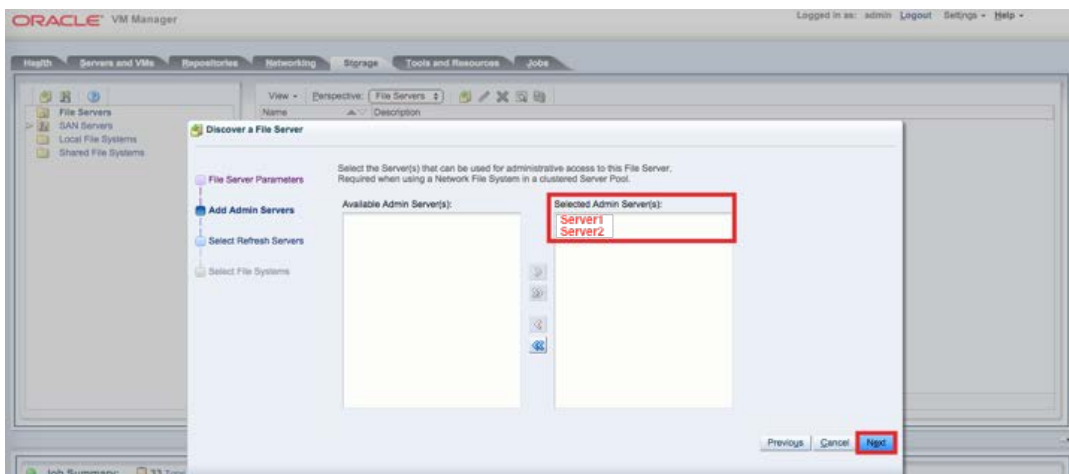


Figure 25. Adding admin servers for Oracle ZFS Storage Appliance in the Oracle VM Manager console



In the next dialog window, corresponding to the highlighted Select Refresh Servers step, select the Oracle VM servers that will have Oracle ZFS Storage Appliance's network file systems exported to them. The servers will discover those network file systems. Click Next.

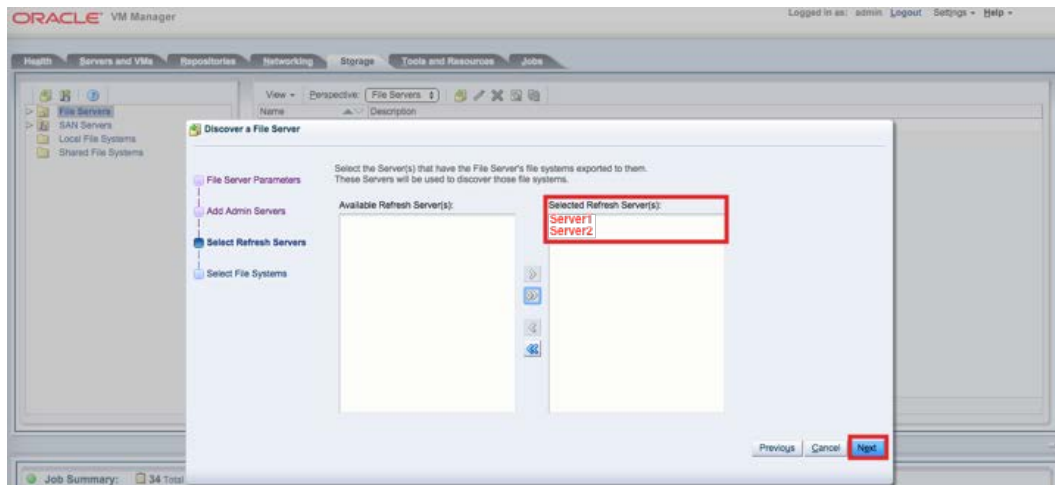


Figure 26. Designating refresh servers for discovering file systems in the Oracle VM Manager console

The final step, as seen in the following figure, is selection of the Oracle ZFS Storage Appliance network file system that will be used as the cluster file system repository for the Oracle VM clustered server pool. Click on the Finish button when done.

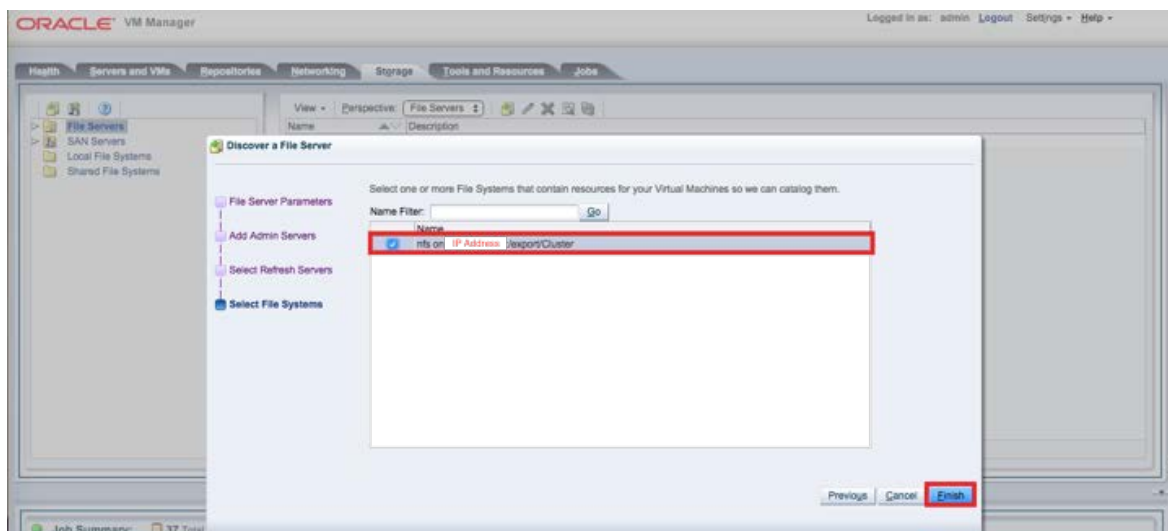


Figure 27. Configuring the Oracle ZFS network file system for the Oracle VM cluster server pool

As seen in the next figure, a new file server – in this case, Oracle ZFS Storage Appliance – will be listed under the File Servers folder, and also listed under Storage, File Servers, and the Perspective list box selection File

Systems. Now the Oracle VM cluster server pool can be configured on top of an Oracle ZFS Storage Appliance network file system.

**IMPORTANT:** Before presenting any Oracle ZFS Storage Appliance network file system to Oracle VM servers, or even before performing any repository configuration on the Oracle VM side, refresh all network file systems. To refresh the network file systems, right click on the network file system which will be used by the Oracle VM repository or presented to the Oracle VM clustered server pool, then click Refresh.

Repeat the same steps for the second head (head 2) of the Oracle ZFS Storage Appliance.

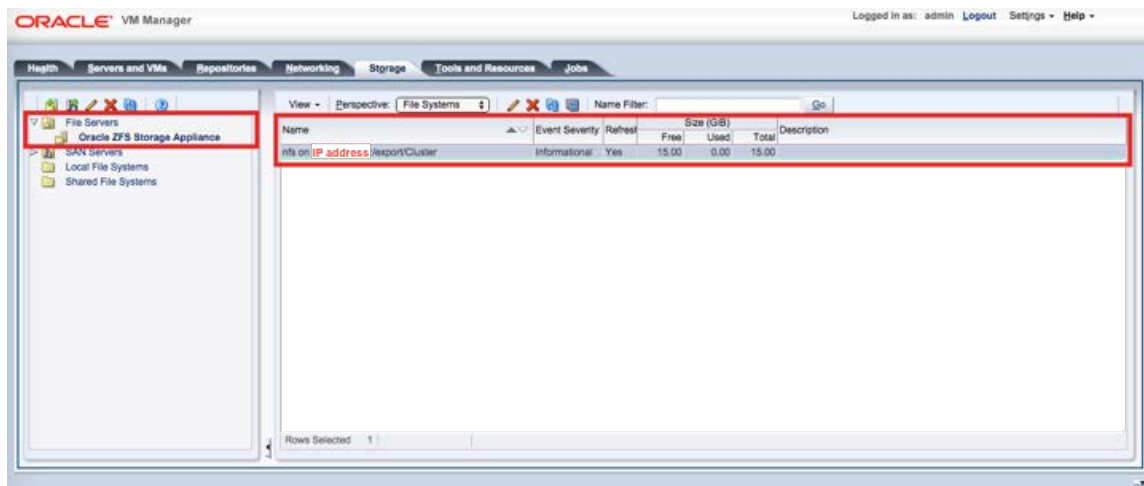


Figure 28. Oracle ZFS Storage Appliance network file system configuration for Oracle VM cluster server pool shown in Oracle VM Manager console

## Configuring the Oracle VM 3 Cluster Server Pool with NFS Protocol

This section presents the best practices for configuring an Oracle VM cluster server pool on top of an Oracle ZFS Storage Appliance network file system.

Assuming that the file server has been successfully presented to the Oracle VM servers, select the Servers and VMs tab in the Oracle VM Manager console, and then click Create Server Pool. In the resulting dialog window as seen in the following figure, provide values for the following settings:

- Server Pool Name – The name of your Oracle VM cluster server pool. In this example, it is Oracle VM Prod-01.
- Virtual IP address for the Pool – The required IP address for the Oracle VM cluster server pool. In this example, it is 10.80.75.25.
- VM Console Keymap – Use en-us (English, United States).
- VM Start Policy – Select Best Server.
- Secure VM Migrate – Leave unchecked. This option can be used for specific cases, which are out of the scope of this document. For further information, please refer to the Oracle VM Administrator's Guide for Release 3.3 at: [http://docs.oracle.com/cd/E50245\\_01/](http://docs.oracle.com/cd/E50245_01/).
- Clustered Server Pool – Check this box.
- Timeout for Cluster – Increase this value to 300 seconds.

- Storage for Server Pool – Select Network File System.
- Storage Locations – Select the Oracle ZFS Storage Appliance network file system previously created (the Cluster network file system).
- Description – Provide a brief description for your server pool.

Click Next.

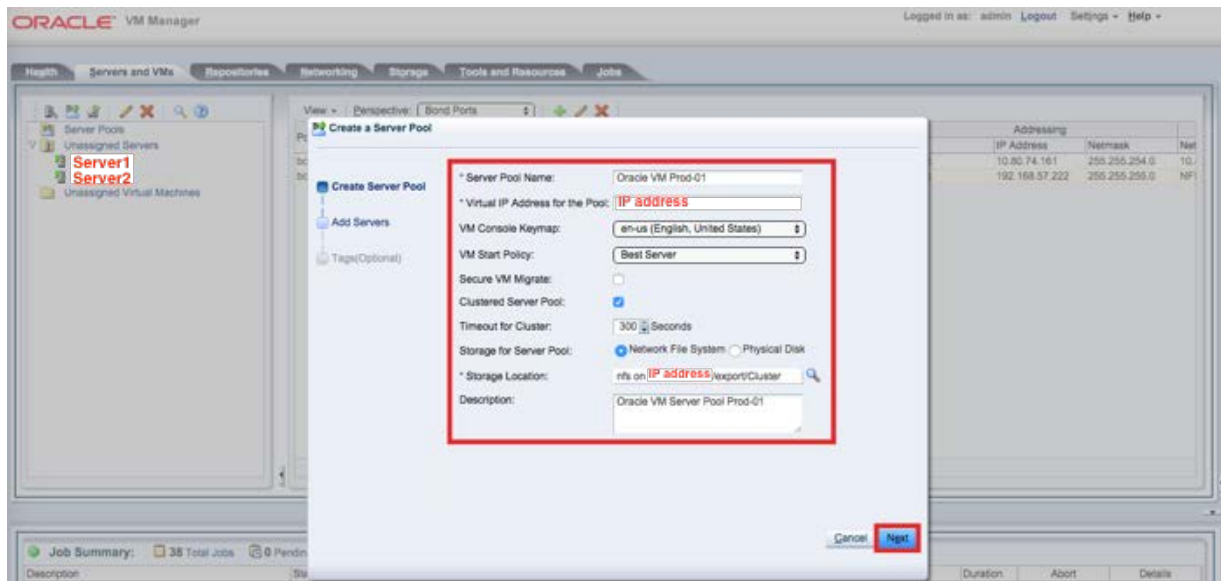


Figure 29. Oracle VM clustered server pool configuration with Oracle ZFS Storage Appliance shown in Oracle VM Manager console

In the next step and associated dialog window, Oracle VM Manager will prompt for servers that will be members of the new Oracle VM clustered server pool. Since the servers utilized in this document's examples are x86 platform, select the pull-down menu option Oracle VM x86 for Hypervisor Filter. Then choose the servers that will be members of the new Oracle VM clustered server pool. Click Move Selected Item to Select Server(s) button, click Next, and then click Finish. The new Oracle VM clustered pool on top of the Oracle ZFS Storage Appliance network file systems will be created.

**IMPORTANT:** Tags configuration is out of the scope of this document, so the tags screen can be ignored. For additional information about Tags, refer to Oracle VM 3.3 documentation.

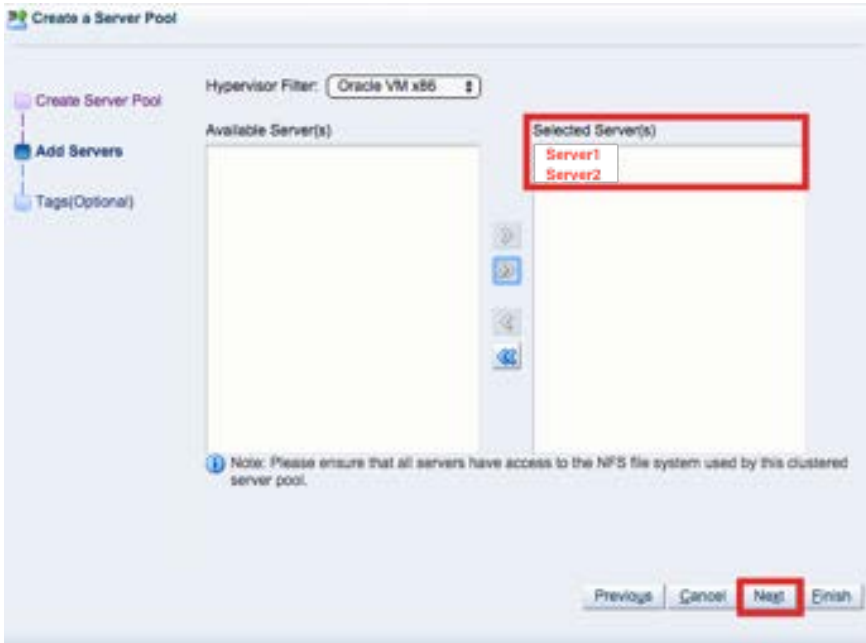


Figure 30. Oracle VM clustered server pool configuration with Oracle ZFS Storage Appliance shown in Oracle VM Manager console

The Oracle VM environment is now ready for new repositories and virtual machine deployments.

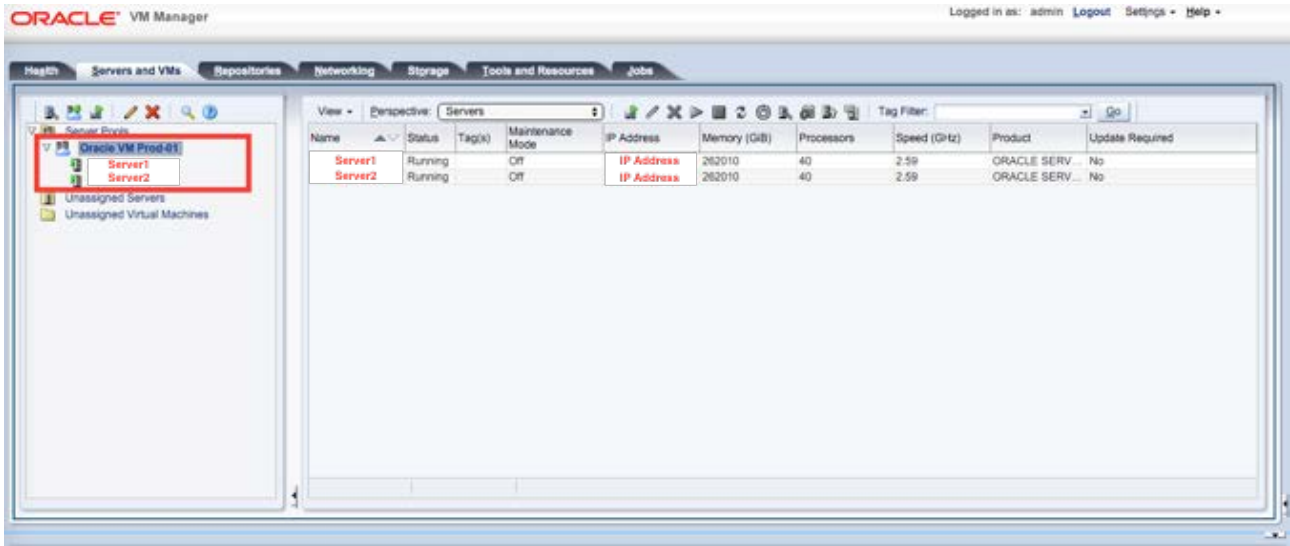


Figure 31. Newly configured Oracle VM clustered server pool

### Configuring Oracle VM 3 Repositories

This section presents the best practices and recommendations for configuring Oracle VM repositories with Oracle ZFS Storage Appliance network file systems.

Be sure to refresh a network file system before configuring it as an Oracle VM repository. After the refresh, navigate in the Oracle VM Manager console to the Repositories tab, click Create a Repository, the green plus (or add) icon on the left hand side of the Oracle VM Manager console, and then enter the following information:

- Repository Name – The name of the new Oracle VM repository. For this example, Database will be the name of the new repository.
- Repository Location – Select Network File Server. Physical disks are chosen only for iSCSI or Fibre Channel (FC) protocol. Both iSCSI and FC protocols are out of the scope of this white paper.
- Network File System – The Oracle ZFS Storage Appliance network file system previously configured and presented to the Oracle VM servers.
- Share Path – Blanked default configuration.
- Description – Brief description for the repository.

Click Next when done.

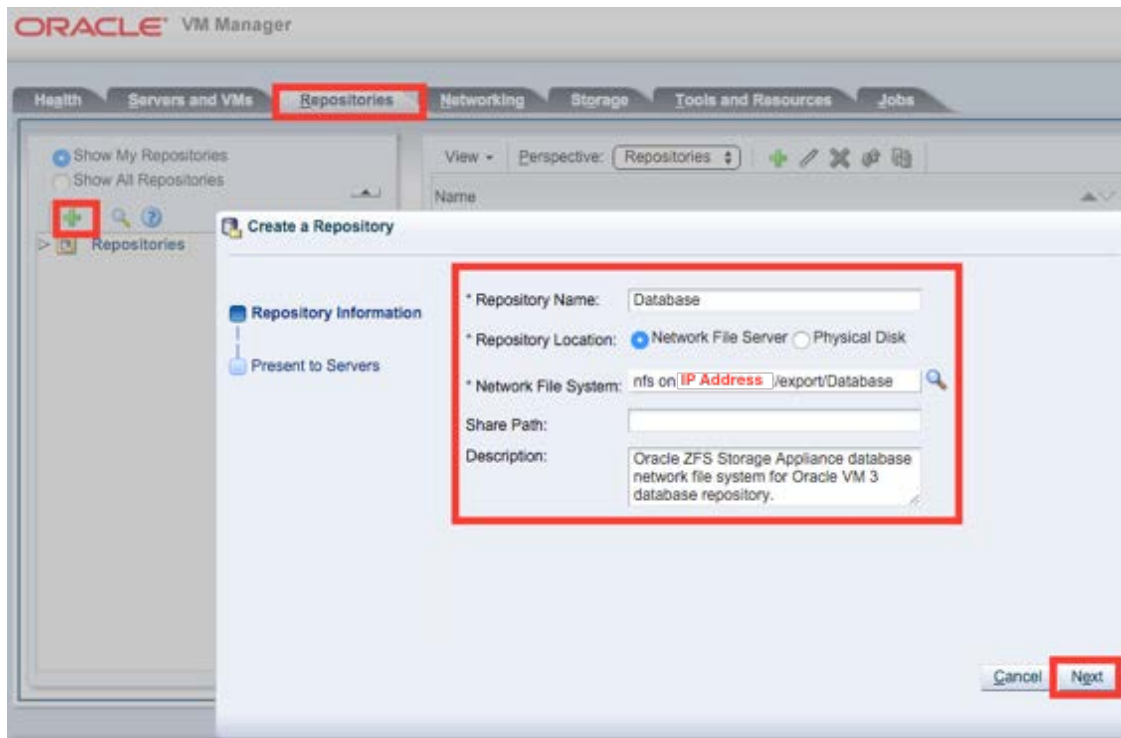


Figure 32. Oracle VM repository information for Oracle ZFS Storage Appliance shown in Oracle VM Manager console

In the next dialog window, select the x86 Oracle VM servers that will receive (be presented with) the new repository. Then click Finish.

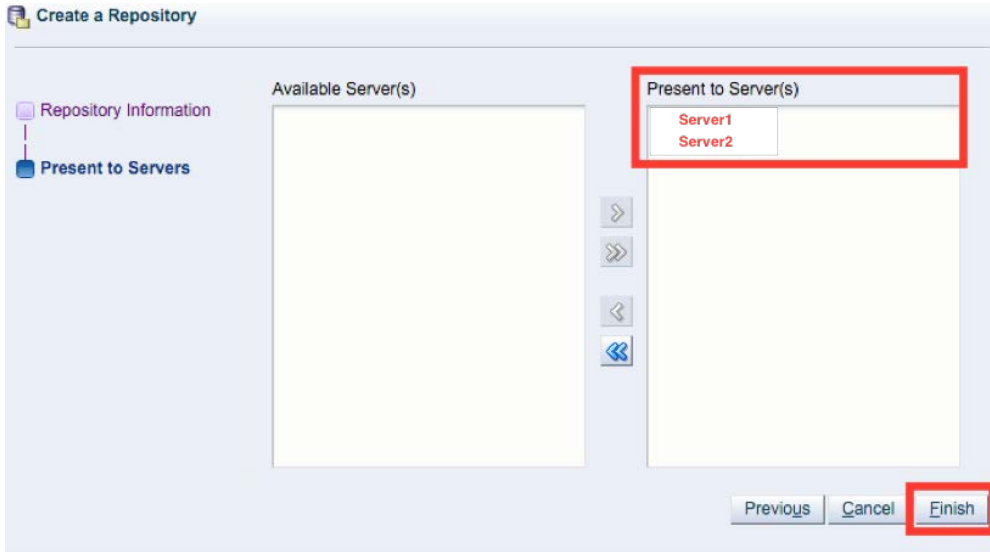


Figure 33. Oracle VM repository configuration shown in Oracle VM Manager console

The new Oracle VM Database repository will be created and available for deployment of assemblies, ISO images, VM files, templates, and virtual machine virtual disks.

The following table presents best practice recommendations for setup of Oracle VM repositories, Oracle ZFS Storage Appliance network file systems, and Oracle ZFS Storage Appliance database record size, as well as the virtual machines and virtual disks that should be deployed in each Oracle VM repository.

**TABLE 1. DATABASE RECORD SIZE RECOMMENDATIONS FOR ORACLE VM REPOSITORIES AND VIRTUAL MACHINES**

Oracle VM Repository	Oracle ZFS Storage Appliance Network File System	Oracle ZFS Storage Appliance Database Record Size	Virtual Machines
Database	/export/Database	32k	Oracle Database, OLTP, Microsoft SQL Servers, MySQL Server virtual disks
E-Business	/export/e-business	128k	E-Business virtual machines
FileServer	/export/FileServer	128k	File Servers virtual machines
MailServer	/export/MailServer	64k	Mail Servers, Microsoft Exchange Servers virtual disks
WebServer	/export/WebServer	128k	Web Servers virtual machines
Windows virtual machine O.S	/export/windows_vms	128k	Windows virtual machines
LINUX virtual machines O.S	/export/linux_vms	128k	Oracle Linux virtual machines

Figure 34 presents all repositories configured and presented to the Oracle VM servers. Each repository is configured on top of a unique Oracle ZFS Storage Appliance network file system.

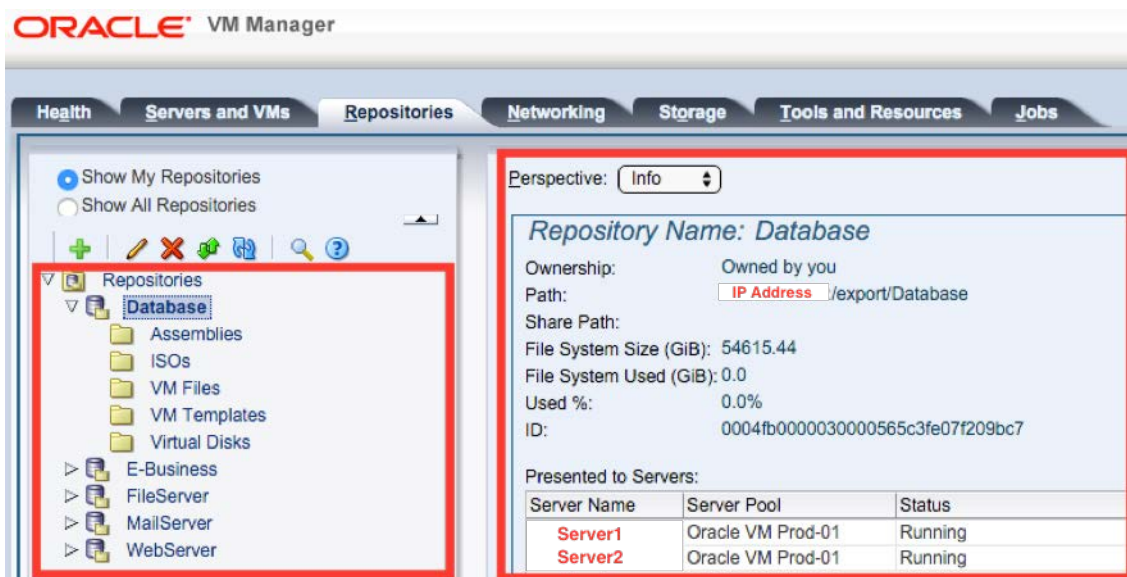


Figure 34. Oracle VM repositories shown in Oracle VM Manager console

## Maximizing Performance with Virtual Machine Setup

Whenever possible, Oracle recommends using paravirtualized virtual machines, as their performance is often superior to the performance of hardware virtualized machines. Paravirtual drivers are optimized and improve the performance of the operating system in a virtual machine. These drivers enable high performance throughput of I/O operations in guest operating systems on top of the Oracle VM Server hosts.

For further information on paravirtual drivers, including installation details for the paravirtual drivers for Microsoft Windows operating systems, see the Oracle Technology Network (OTN) web page "Oracle VM Windows Paravirtual (PD) Drivers at: <http://www.oracle.com/us/technologies/virtualization/virtualization-066470.html>

For additional information related to performance among Oracle VM guest virtualization modes such as Hardware Virtualized (HVM), Paravirtualized (PVM), and Hardware Virtualized with Paravirtualized Drivers (PVHVM), refer to the Oracle Support Document 757719.1 – Comparison of Guest Virtualization Modes; HVM, PVM and PVHVM. Configuration, Mode Validation and Conversion to PVHVM on the My Oracle Support (MOS) web site listed at the end of this document in References.

## Virtual Disk Tuning for Oracle ZFS Storage Appliance

As part of the best practices and recommendation, increasing virtual disk timeout of the virtual machines' virtual disks is critical. A higher SCSI timeout value will ensure that the virtual machine and its applications survive during longer storage failback and takeover operations. In many cases, and also depending of a number of factors, storage failback and takeover operations can take longer than expected – usually longer than 60 seconds. By default, Windows Server operating systems releases 2008, 2008R2, 2012, 2012R2, Windows 7 and Windows 8 have the default SCSI timeout set to 60 seconds, which can be problematic if for some reason the storage



fallback or takeover operations take longer than that. To avoid this problem, increase the SCSI disk timeout values in the Windows virtual machine guest operating systems.

To change the SCSI disk timeout parameter in a Windows virtual machine, follow these instructions:

1. Open the **REGEDIT** (the Registry Editor) of the system, then use the navigation tree to go to **HKEY\_LOCAL\_MACHINE/System/CurrentControlSet/Services/Disk**.
2. Find the option called **TimeOutValue**. Right-click this option, then select **Modify**.
3. Select **Decimal Value**, increase the field **Value data** to **300**, then click **OK**.
4. Reboot the virtual machine.

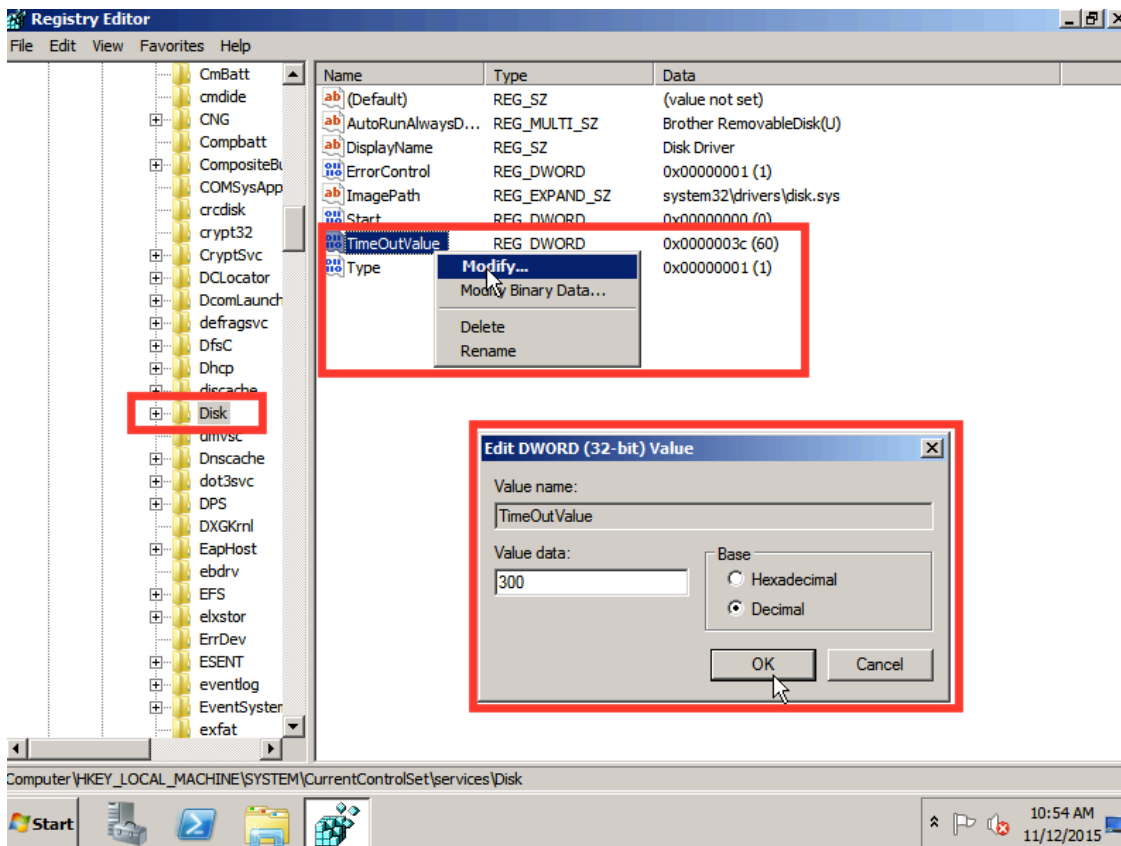


Figure 35. Windows 2008R2 virtual machine REGEDIT shown in the Oracle VM Manager console

## Overview of Tested System Components

The following tables describe the hardware configuration, operating systems, and software releases utilized by systems under test for findings highlighted in this white paper.

**TABLE 2. HARDWARE USED BY THIS WHITE PAPER**

Equipment	Quantity	Configuration
Storage	1 cluster (2 controllers)	Oracle ZFS Storage ZS3-2 cluster 256 GB DRAM per controller 2 x 512 GB read cache SSD per controller One x 24 3-terabyte (TB) disk drivers – Oracle Storage Drive Enclosure DE2-24C One x 20 3-terabyte (TB) disk drivers – Oracle Storage Drive Enclosure DE2-24C One x 24 900-gigabyte (GB) disk drivers – Oracle Storage Drive Enclosure DE2-24P One x 20 900-gigabyte (GB) disk drivers – Oracle Storage Drive Enclosure DE2-24P Four x 10 gigabit Ethernet (GbE) network interface cards (NICs) per controller Four x 73 GB log devices
IP Network Switch	2	10 GbE network switch
X86_64 Servers	3	X86_64 for Oracle VM and Oracle VM Manager servers

**TABLE 3. VIRTUAL MACHINE COMPONENTS USED BY THIS DOCUMENT**

Operating System	Quantity	Configuration
Oracle Linux 6	2	Oracle Linux 6.7 x86_64 virtual machines
Oracle Linux 7	2	Oracle Linux 7.1 x86_64 virtual machines
Microsoft Windows 2008R2	2	Microsoft Windows 2008R2 x86_64 virtual machines
Microsoft Windows 2012R2	2	Microsoft Windows 2012R2 x86_64 virtual machines

**TABLE 4. SOFTWARE RELEASES USED BY THIS WHITE PAPER**

Software	Version
Oracle ZFS Storage Appliance Software	OS 8.4.2
Oracle VM Servers	3.3.3-1085
Oracle VM Manager	3.3.3-1085



## Conclusion

The outstanding performance of Oracle ZFS Storage Appliance, coupled with Oracle VM's facile management for deploying Oracle Database and numerous other Oracle applications in virtualized environments, provides superior overall performance and efficiency that can reduce your total cost of ownership. This makes them ideal choices for your virtual environment.

Implementing configuration best practices and recommendations, tailored to your business needs, for these products will help ensure that you maximize the inherent high performance they offer.



## References

See the following resources for additional information relating to the products covered in this document.

- Oracle ZFS Storage Appliance White Papers and Subject-Specific Resources  
<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html>
- Oracle ZFS Storage Appliance Product Information  
<https://www.oracle.com/storage/nas/index.html>
- Oracle ZFS Storage Appliance Documentation Library, including Installation, Analytics, Customer Service, and Administration guides:  
<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>
- The *Oracle ZFS Storage Appliance Administration Guide* is also available through the Oracle ZFS Storage Appliance help context.  
The Help function in Oracle ZFS Storage Appliance can be accessed through the browser user interface.
- "Oracle VM Windows Paravirtual (PD) Drivers" – information from Oracle Technology Network's Oracle Virtualization information site  
<http://www.oracle.com/us/technologies/virtualization/virtualization-066470.html>



**Oracle Corporation, World Headquarters**

500 Oracle Parkway  
Redwood Shores, CA 94065, USA

**Worldwide Inquiries**

Phone: +1.650.506.7000  
Fax: +1.650.506.7200

---

CONNECT WITH US

-  [blogs.oracle.com/oracle](http://blogs.oracle.com/oracle)
-  [facebook.com/oracle](http://facebook.com/oracle)
-  [twitter.com/oracle](http://twitter.com/oracle)
-  [oracle.com](http://oracle.com)

## Integrated Cloud Applications & Platform Services

Copyright © 2015, Oracle and/or its affiliates. All rights reserved. This document is provided *for* information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0615

Oracle VM 3 with Oracle ZFS Storage Appliance: Configuration Best Practices for Oracle ZFS Storage Appliance  
Using NFS v4 Protocol  
December 2015  
Author: Anderson Souza, Application Integration Engineering



Oracle is committed to developing practices and products that help protect the environment.