

# Best Practices for Configuring Oracle ZFS Storage Appliance and VMware vSphere 6.x with NFS Protocol

ORACLE WHITE PAPER | APRIL 2017





## Table of Contents

Introduction	1
About Oracle ZFS Storage Appliance Systems	2
Best Practices for Deploying VMware vSphere 6.x, Oracle ZFS Storage Appliance Systems, and the NFS Protocol	2
Controllers, Software Release, and Disk Pools	2
Oracle ZFS Storage Appliance CPU, L1, and L2 Cache	3
Network Settings	4
NFS, Projects, and Shares	5
VMware vSphere 6.x Host Network Infrastructure	8
Recommendations for the NFS Protocol	14
VMware vSphere 6.x Cluster Recommendations	15
Virtual Machine and Data Layout Recommendations	16
Monitoring VMware vSphere 6.x with DTrace Analytics and the esxtop Tool	17
Monitoring NFS Performance	17
Conclusion	19
Appendix A: Benchmark Results	20
Appendix B: References	20



## Introduction

This white paper provides best practices and recommendations for configuring VMware vSphere 6.x and the Oracle ZFS Storage Appliance family of products with the NFS protocol to reach optimal I/O performance and throughput.

The outlined best practices and recommendations highlight configuration and tuning options as well as recommendations for the correct design of VMware vSphere 6.x with the NFS protocol and Oracle ZFS Storage Appliance systems.

Highlighted in this paper are the following:

- » Best practices and recommendations for employing VMware vSphere 6.x with Oracle ZFS Storage Appliance systems
- » Tuning options for the NFS protocol in production environments
- » IP network design for NFS datastores with Oracle ZFS Storage Appliance systems
- » VMware virtual machine (VM) data layout
- » Monitoring options using VMware's esxtop tool and the Oracle ZFS Storage Appliance DTrace Analytics feature

This document is intended for experienced virtualization administrators, system engineers, and systems architects working with VMware products and Oracle ZFS Storage Appliance. To benefit from its information requires basic knowledge of networking, NAS storage, and Microsoft Windows operating systems.

**Note:** References to Sun ZFS Storage Appliance, Sun ZFS Storage 7000, and ZFS Storage Appliance all refer to the same Oracle ZFS Storage Appliance family.

## About Oracle ZFS Storage Appliance Systems

The basic architecture of the products in the Oracle ZFS Storage Appliance family is designed to provide high performance, flexibility, and scalability. Oracle ZFS Storage Appliance systems provide multiple connectivity protocols for data access, including Network File System (NFS), Common Internet File System (CIFS), Internet Small Computer System (iSCSI), InfiniBand (IB), and Fibre Channel (FC). It also supports the Network Data Management Protocol (NDMP) for backing up and restoring data. The Oracle ZFS Storage Appliance architecture also offers the Hybrid Storage Pool feature, in which direct random access memory (DRAM), flash, and physical disks are seamlessly integrated for efficient data placement. A powerful performance monitoring tool called DTrace Analytics provides details about the performance of various components, including network, storage, file systems, and client access. The tool also offers plenty of drill-down options that allow administrators to monitor specific rates of latency, size of transfers, and utilization of resources. Oracle ZFS Storage Appliance systems provide a variety of RAID protocols to balance the capacity, protection, and performance requirements of applications, databases, and virtualized environments.

## Best Practices for Deploying VMware vSphere 6.x, Oracle ZFS Storage Appliance Systems, and the NFS Protocol

The following configurations for Oracle ZFS Storage Appliance systems are recommended to optimize performance with VMware vSphere 6.x.

### Controllers, Software Release, and Disk Pools

Virtualized environments produce high amounts of random I/O patterns and need high storage performance as well as high availability, low latency, and fast response time. To meet these demands, use a mirrored data profile. This configuration duplicates copies as well as produces fast and reliable storage by dividing access and redundancy—usually between two sets of disks. In combination with write solid-state drive (SSD) log devices and the Oracle ZFS Storage Appliance architecture, this profile can produce a large amount of input/output operations per second (IOPS) to attend to the demand of critical virtual desktop environments.

The recommended minimum disk storage configuration for VMware vSphere 6.x includes

- » A mirrored disk pool with the following configuration:
  - » For models using high-performance (HP) disks: Use at least twenty 600 GB 10,000 RPM HP 2.5-inch SAS3 hard disk drives (HDDs)—or at least twenty 1.2 TB 10,000 RPM HP 2.5-inch SAS3 HDDs—and at least two 200 GB SSD devices for LogZilla working with a striped log profile.
  - » For models using high-capacity (HC) disks: Use at least forty-four 8 TB 7,200 RPM HC 3.5-inch SAS3 HDDs and at least two 200 GB SSD devices for LogZilla working with a striped log profile.
- » At least two 1.6 TB SSDs for Level 2 Adaptive Replacement Cache (L2ARC): Use a striped cache for both the HC and HP models.

**Note:** The following example demonstrates an Oracle ZFS Storage Appliance system with two Oracle Storage Drive Enclosure DE2-24C trays fully populated with forty-four 3 TB SAS2 7,200 RPM disks and two 1.6 TB SSDs devices for ReadZilla configured in striped mode, and two 200 GB SSD devices configured for LogZilla, also in striped mode. See Figure 1 and Figure 2.

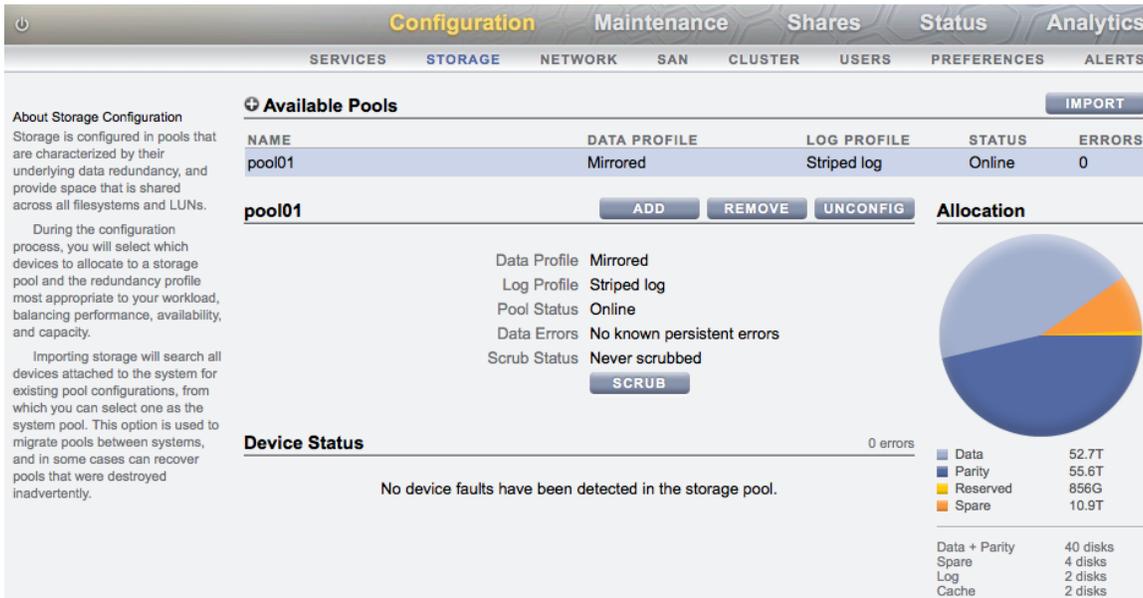


Figure 1. Disk pool configuration in the Oracle ZFS Storage Appliance browser user interface (BUI)

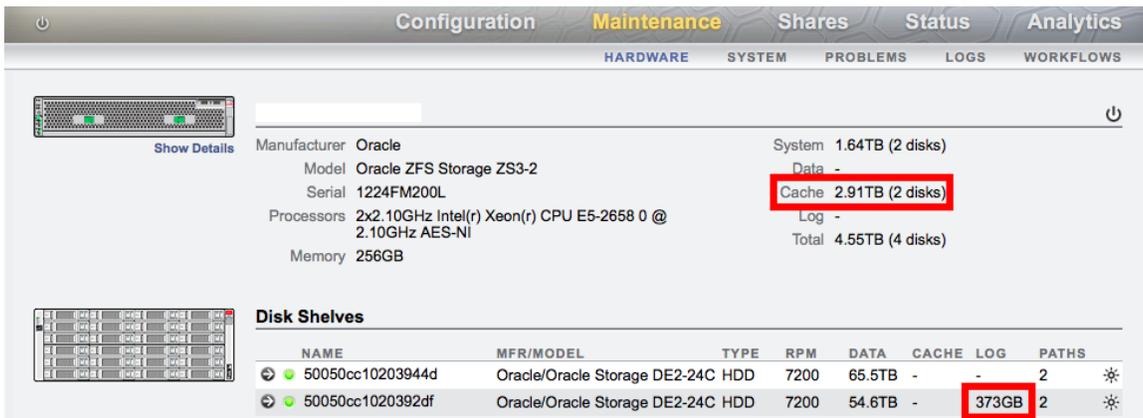


Figure 2. Log devices (LogZilla) and cache devices (ReadZilla) in the Oracle ZFS Storage Appliance BUI

For high availability and proper load balancing for a VMware vSphere 6.x environment with the NFS protocol, use an Oracle ZFS Storage Appliance model that supports clustering. Configure the cluster in active/active mode and use Oracle ZFS Storage Appliance software release OS8.7 (2013.1.7.x) or greater. Oracle ZFS Storage Appliance software releases can be downloaded from My Oracle Support. Look for "Oracle ZFS Storage Appliance: Software Updates" (Doc ID 2021771.1). Refer to "Appendix B: References" of this document for further information.

### Oracle ZFS Storage Appliance CPU, L1, and L2 Cache

The following combination and sizing of CPU, Level 1 Adaptive Replacement Cache (L1ARC), and L2ARC are critical to meet the demands of compression algorithms and operations as well as to meet overall performance in large deployments of VMware vSphere 6.x with Oracle ZFS Storage Appliance systems.



The minimum recommend configuration is

- » At least two 2 GHz Intel Xeon CPUs (Intel® Xeon® processor X7750 at 2.00 GHz) per Oracle ZFS Storage Appliance head.
- » At least 512 GB of DRAM memory (L1ARC) per Oracle ZFS Storage Appliance head.
- » Because most of the reading operations for VMware workloads will be provided primarily by Oracle ZFS Storage Appliance DRAM/ARC (L1ARC), ReadZilla SSD devices are not mandatory, but they are recommended. So, as part of the Oracle ZFS Storage Appliance architecture, at least two 1.6 TB SSD devices are recommended for VMware environments.

## Network Settings

The following best practices for network configuration are designed to ensure that NFS traffic between Oracle ZFS Storage Appliance systems and VMware vSphere 6.x hosts achieves high availability and no single point of failure:

- » Use NFSv4.1 protocol with Oracle ZFS Storage Appliance systems.
- » Isolate the storage traffic from other networking traffic. You can configure this utilizing VLAN, network segmentation, or dedicated switches for NFS traffic between Oracle ZFS Storage Appliance systems and VMware vSphere 6.x hosts only.
- » On Oracle ZFS Storage Appliance systems, configure at least two physical 10 GbE (dual-port) network interface cards (NICs) per head, bundled into a single channel using IEEE 802.3ad Link Aggregation Control Protocol (LACP) with a large maximum transmission unit (MTU) jumbo frame (9,000 bytes). If you are working with a cluster configuration, configure at least two 10 GbE (dual-port) NICs per Oracle ZFS Storage Appliance head, and also use an IP network multipathing (IPMP) configuration in combination with LACP. With an IPMP configuration, you will achieve network high availability, and with link aggregation you will obtain better network performance. These two technologies complement each other and can be deployed together to provide benefits for network performance and availability for VMware vSphere 6.x with Oracle ZFS Storage Appliance systems.
- » For picking an outbound port based on the source and IP address, utilize LACP policy L3.
- » For switch communication mode, use the LACP active mode, which will send and receive LACP messages to negotiate connections and monitor the network link status.
- » Use an LACP “Short” timer interval between LACP messages, as seen in the configuration in Figure 3.

**Note:** Some IP network switch vendors do not support the LACP protocol. In this situation, set the LACP mode to “Off.” Please refer to your IP switch vendor documentation for more information.

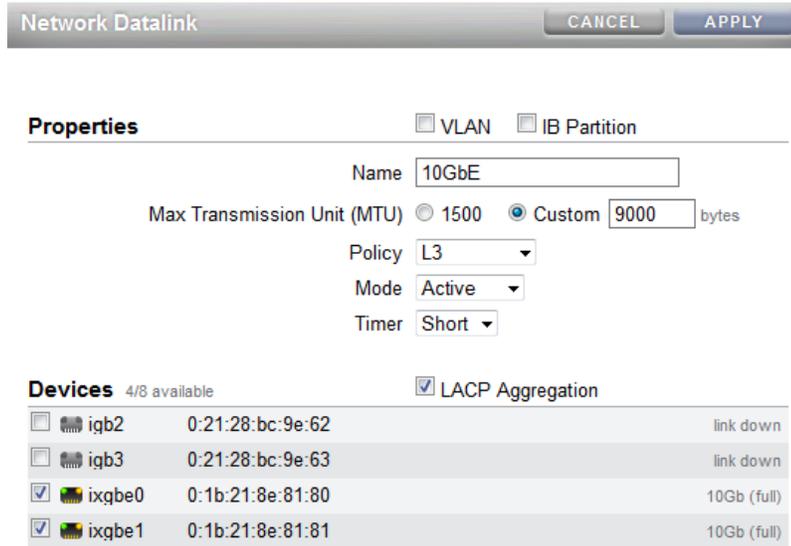


Figure 3. LACP, jumbo frame, and MTU configurations in the Oracle ZFS Storage Appliance BUI

## NFS, Projects, and Shares

When working with an Oracle ZFS Storage Appliance system with more than one disk shelf, try to split the workload across different disk pools and use the “No Single Point of Failure” (NSPF) feature. This design will provide you with more storage resources as well as better I/O load balancing, performance, and throughput for your VMware environments.

Oracle ZFS Storage Appliance software release OS8.7 provides support to NFS protocol version 4.1 (NFSv4.1). The following recommendations are best practices for Oracle ZFS Storage Appliance NFSv4.1 with VMware vSphere 6.x:

- » To enable NFSv4.1 in the Oracle ZFS Storage Appliance BUI, select **Configuration > Services > NFS**. On the NFS configuration screen, ensure that **Minimum supported version** is set to NFSv3, and **Maximum supported version** is set to NFSv4.1.
- » Ensure that the maximum number of server threads is set to at least 3,000 and the grace period is set to 90 seconds.
- » Ensure that the **Enable NFSv4 delegation** option is selected and that **Mount visibility** is set to Full.
- » All other options, such as **Custom NFSv4 numeric id domain** and **Use NFSv4 numeric id strings**, should be set to the default (that is, disabled).

Figure 4 shows the NFS service configuration on the Oracle ZFS Storage Appliance BUI.

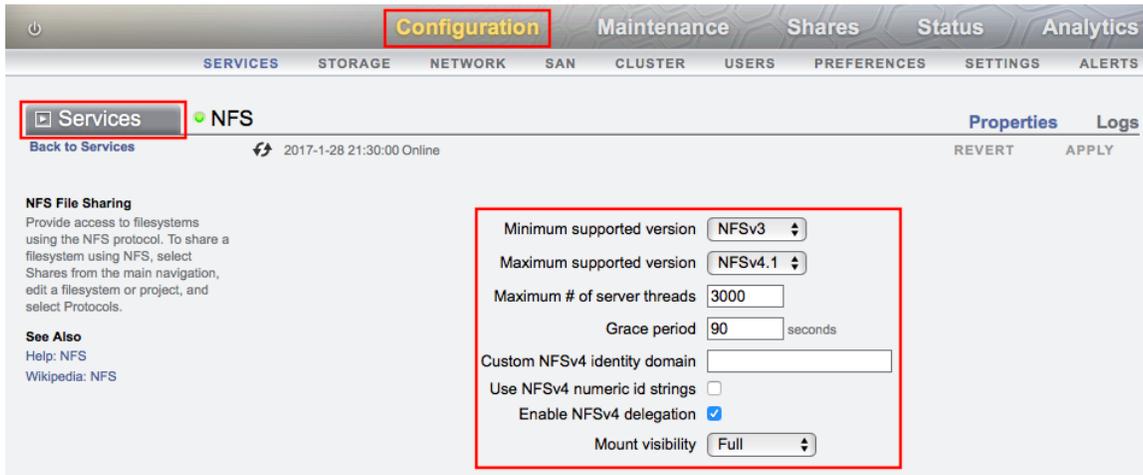


Figure 4. LACP, jumbo frame, and MTU configurations in the Oracle ZFS Storage Appliance BUI

For higher performance and throughput, some applications issue large I/O requests to the storage devices. By default, the vmkernel issues I/O requests as large as 32,767 KB (32 MB). So ensure that the `Disk.DiskMaxIOSize` advanced option of the VMware vSphere 6.x host is configured to at least 32,767 KB (32 MB). This setting should prevent performance bottlenecks on the ESXi host's side and provide for better performance of VMware vSphere 6.x with Oracle ZFS Storage Appliance systems.

**Note:** Adjustments to the `Disk.DiskMaxIOSize` parameter need to be carefully evaluated. For most use cases, it is better to leave this parameter at its default (32,767 KB/32 MB) block size for increased performance and lower CPU utilization on the ESXi host's side.

For additional information regarding `Disk.DiskMaxIOSize` advanced settings on ESXi hosts, refer to VMware Knowledge Base article "Tuning ESXi for better storage performance by modifying the maximum I/O block size" (VMware KB 10034690) listed in the "Appendix B: References" section of this white paper.

Table 1 lists Oracle ZFS Storage Appliance projects and file systems along with the recommended database record size per application.

**TABLE 1: POOL'S PROJECTS AND FILE SYSTEMS SHARES CREATED FOR VMWARE VSPHERE 6.X PRODUCTION ENVIRONMENTS**

Project	File Systems	Recommended Database Record Size
VM_OS (Virtual machines' operating systems or virtual machine disks [VMDKs])	/export/vm_os	32 KB: VMkernel issues I/O operations with 32 KB block size
EXCHANGE_DB (Microsoft Exchange database (DB) VMDKs)	/export/ms-exchangedb	32 KB for Microsoft Exchange Database and logs
Microsoft SQL DB VMDKs (Microsoft SQL Server database VMDKs) <b>Note:</b> Due to different block sizes, Microsoft SQL DB can have at least three (or more) different Oracle ZFS Storage Appliance file systems.	/export/mssqldb export/mssqldb_data export/mssqldb_backup	OLTP LOG transactions: Can be any size of up to 64 KB, which provides the best performance. OLTP DATA operation: 8 KB. Bulk insert operations: Can be any multiple of 8 KB up to 256 KB. Read Ahead (DSS index scans): Can be a multiple of 8 KB up to 512 KB. Microsoft SQL Backup operation: 1 MB.
File Server/Data VMDKs	/export/fsdata	32 KB, 128 KB. or 1 MB

Figure 5 shows the file systems and mountpoint configurations in the Oracle ZFS Storage Appliance BUI. Details for the configuration choices and recommendations follow.

**Space Usage**

**DATA**

Quota  0 G ▾  
 Include snapshots

Reservation  0 G ▾  
 Include snapshots

**USERS & GROUPS**

User or Group ▾   
 Show All

Usage none  
 Quota  None  Default  
 0 G ▾

**Properties**

Inherit from project

Mountpoint

Read only

Update access time on read

Non-blocking mandatory locking

Data deduplication (warning)

Data compression  ▾

Checksum  ▾

Cache device usage  ▾

Synchronous write bias  ▾

Database record size  ▾

Additional replication  ▾

Virus scan

Prevent destruction

Restrict ownership change

Figure 5. Share configuration for VMware vSphere 6.x shown in Oracle ZFS Storage Appliance BUI

- » **Read only** option: Leave disabled.
- » **Update access time on read**: Disable this option. This option is valid only for file systems, and controls whether the access time for files is updated upon read. Under heavy workloads consisting primarily of reads, and over large number of files, turning this option off might improve performance for VMware virtual machines.
- » **Non-blocking mandatory locking**: Do not select this option. This option is valid only for file systems for which the primary protocol is Server Message Block (SMB). SMB is not covered by this white paper, and it is not part of the VMware environment with Oracle ZFS Storage Appliance systems.
- » **Data deduplication (warning)** option: Do not select this option.  
**Note:** Data deduplication is supported only with Oracle ZFS Storage Appliance software release OS8.7 and greater. Also, this feature is recommended and supported only for backup use cases, and it is not intended to be used with VMware virtual machines in production. Do not enable data deduplication on Oracle ZFS Storage Appliance file systems on which VMware virtual machines have been deployed and are in production.
- » **Data compression**: Select the LZ4 algorithm as data compression. Before writing data to the storage pool, shares can optionally compress data utilizing different algorithms of compression. The LZ4 algorithm is considered a fast compression algorithm and it does not consume many CPU resources. For data reduction purposes and disk space efficiency, it is recommended to use the LZ4 compression algorithm with VMware virtual machines and Oracle ZFS Storage Appliance systems.
- » **Checksum**: Select the Fletcher4 (Standard) checksum algorithm. This feature controls the checksum algorithm used for data blocks, and it also allows the system to detect invalid data returned from devices. Working with the Fletcher4 algorithm is sufficient for normal operations and can help avoid additional CPU load.
- » **Cache device usage**: The “All data and metadata” option is recommended. With this option, all virtual machine files and any metadata generated by the virtual machines will be cached by the Oracle ZFS Storage Appliance system.
- » **Synchronous write bias**: To provide faster response time for VMware, select the “Latency” option.
- » **Database record size**: Configure this setting according to Table 1.
- » **Additional replication**: To store a single copy of data blocks, select the “Normal (Single Copy)” option.
- » **Virus scan**: Assuming that each virtual machine will have its own antivirus software up and running, enabling virus scan here is not recommended. However, in a virtualized environment, this option can be enabled for an additional NFS share that hosts a Windows home directory or shared folder for all users.  
For additional information regarding virus scan solutions with Oracle ZFS Storage Appliance systems—for example, McAfee VirusScan Enterprise solution with Oracle ZFS Storage Appliance systems—refer to the Oracle Technology Network link provided in “Appendix B: References,” where you can find supporting white papers and documentation for Oracle ZFS Storage Appliance system, including the “How to Configure McAfee VirusScan Enterprise for the Oracle ZFS Storage Appliance” white paper.”
- » **Prevent destruction**: By default, this option is turned off. Enabling this option to protect the NFS share from accidental destruction is recommended.
- » **Restrict ownership change**: By default, this option is turned on. Also, for this example, changing the ownership of virtual machine files is not recommended.

## VMware vSphere 6.x Host Network Infrastructure

The example that follows employs the recommended network configuration for VMware vSphere 6.x hosts with an Oracle ZFS Storage Appliance system:

- » There are at least two 10 GbE NICs connected to different 10 GbE IP switches with all interfaces working with 10,000 Mb, full duplex, and 9,000 MTU (jumbo frame).
- » Ensure that the VMware VMkernel port adapter associated with this virtual switch has the VMware vSphere vMotion traffic option enabled. This configuration will enable VMware vSphere vMotion traffic and operations on these physical interfaces. See Figure 6.

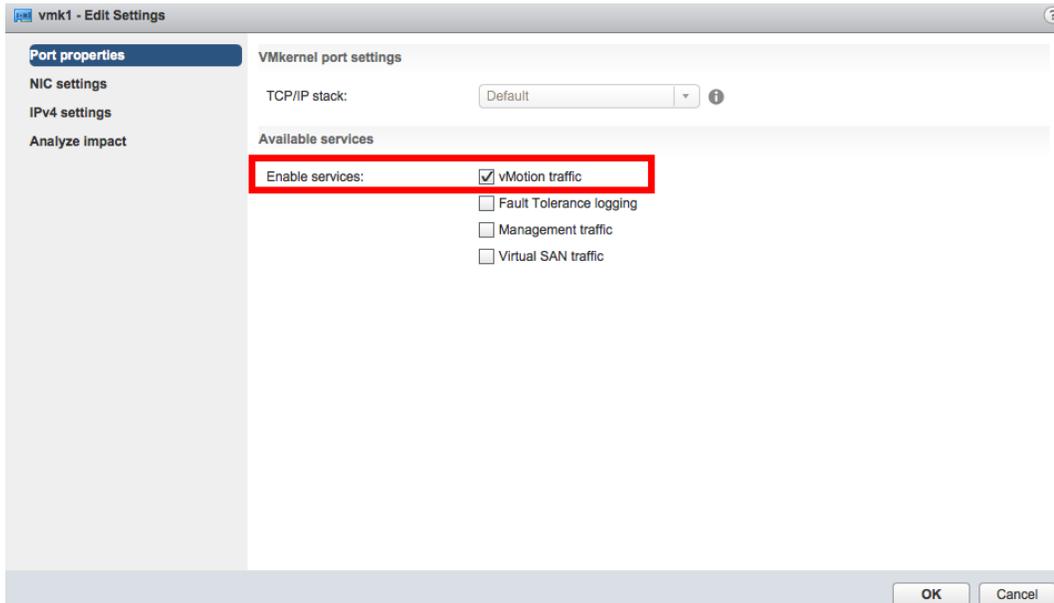


Figure 6. VMware vSphere vMotion configuration shown on VMware vSphere VCenter Server Web Client

» The VMware SR-IOV (Single Root I/O Virtualization) feature is also supported by Oracle ZFS Storage Appliance systems with VMware vSphere 6.x hosts, and it can be enabled as well. Assuming that the VMware vSphere host is using an SR-IOV-capable network card, individual virtual functions (VFs) on the physical network device can be assigned to VMware virtual machines in pass-through mode (VMDirectPath I/O mode). SR-IOV is recommended especially for workloads or critical tier 1 applications that require low-latency network characteristics. See Figure 7.

**Note:** For additional information related to SR-IOV, refer to “vSphere Networking,” which is listed in “Appendix B: References.”

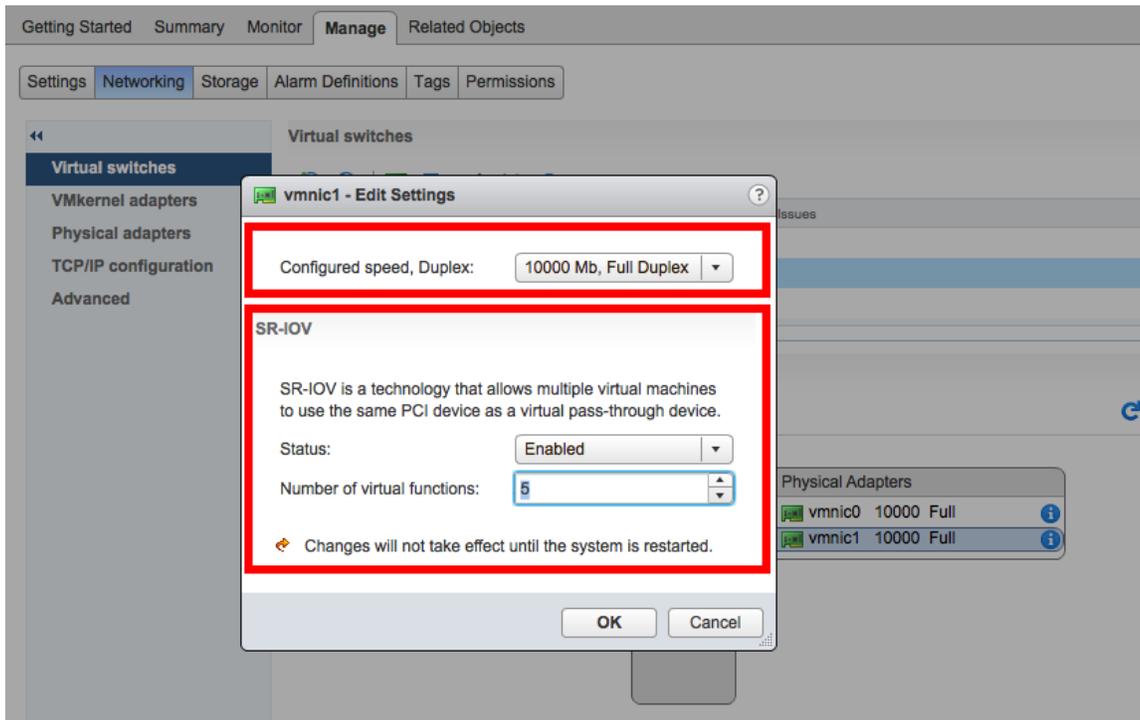


Figure 7. VMware vSphere physical adapter speed, duplex, and SR-IOV configurations shown on VMware vSphere VCenter Server Web Client

- » Network security policies allow VMware administrators to define network filters based on inbound and outbound network frames. As part of the best practices and recommendations for VMware vSphere 6.x production environments with an Oracle ZFS Storage Appliance system and the NFS protocol, select the following properties, as seen in Figure 8:
  - » **Promiscuous mode:** Set to “Reject.”
  - » **MAC address changes:** Set to “Accept.”
  - » **Forged transmits:** Set to “Accept.”

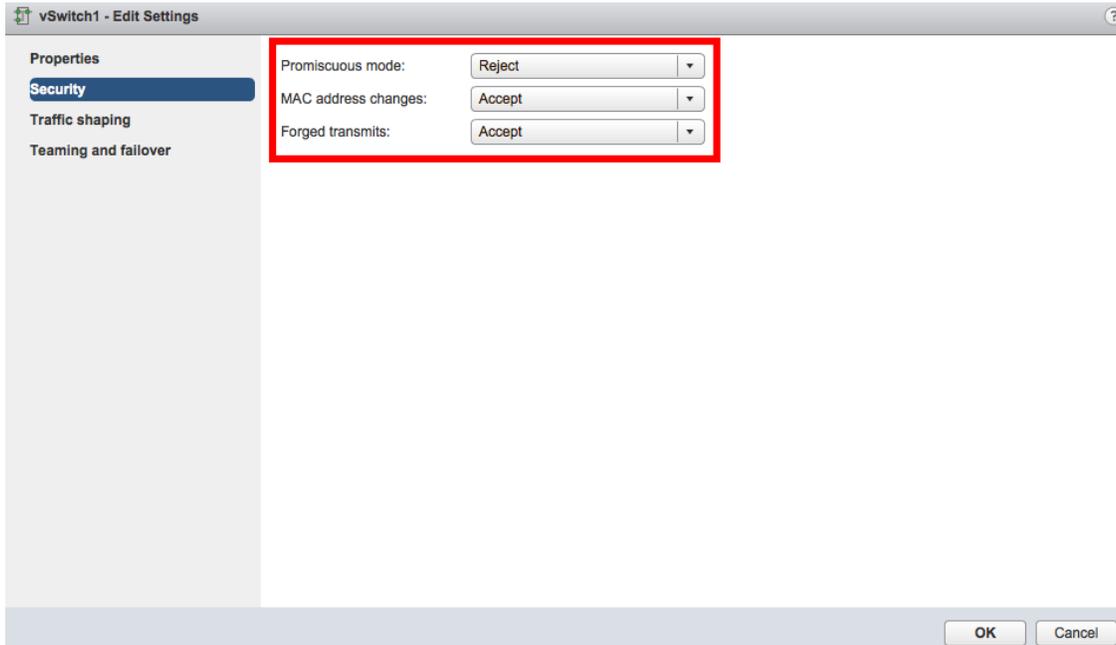


Figure 8. VMware vSphere virtual switch security configuration shown on VMware vSphere VCenter Server Web Client

- » Traffic shaping allows VMware administrators to shape the network outbound traffic on VMware vSphere standard switches. This policy is used to restrict the network bandwidth available to any vSphere host port, and it is defined based on three different characteristics: average bandwidth, peak bandwidth, and burst size. For VMware vSphere 6.x production environments with an Oracle ZFS Storage Appliance system and the NFS protocol, set the **Status** option in the traffic shaping configuration of the VMware virtual switch to “Disabled.” See Figure 9.

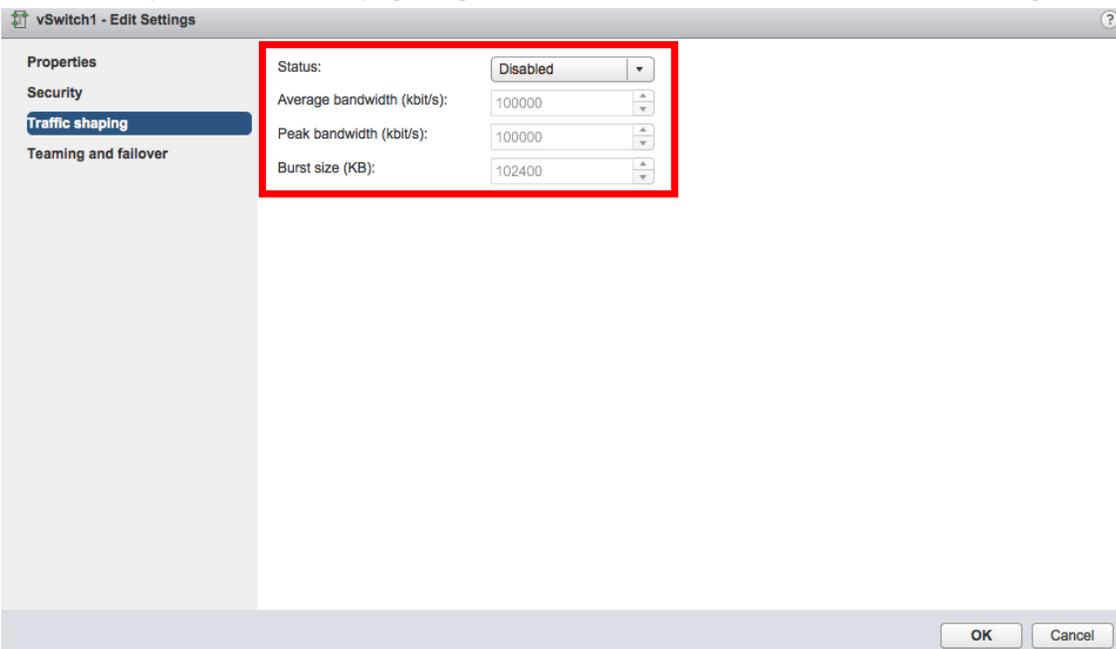


Figure 9. VMware vSphere virtual switch traffic shaping configuration shown on VMware vSphere VCenter Server Web Client

» The **Active adapters** option of the VMware vSphere virtual switch configuration should be set to active/active mode. See Figure 10.

On the VMware side and for at least two or more NICs that are a member of a port-channel group, use the VMware NIC teaming configuration shown below and in Figure 10.

- » **Load balancing:** Select “Route based on IP hash.”
- » **Network failure detection:** Select “Link status only.”
- » **Notify switches:** Select “Yes.”
- » **Failback:** Select “Yes.”

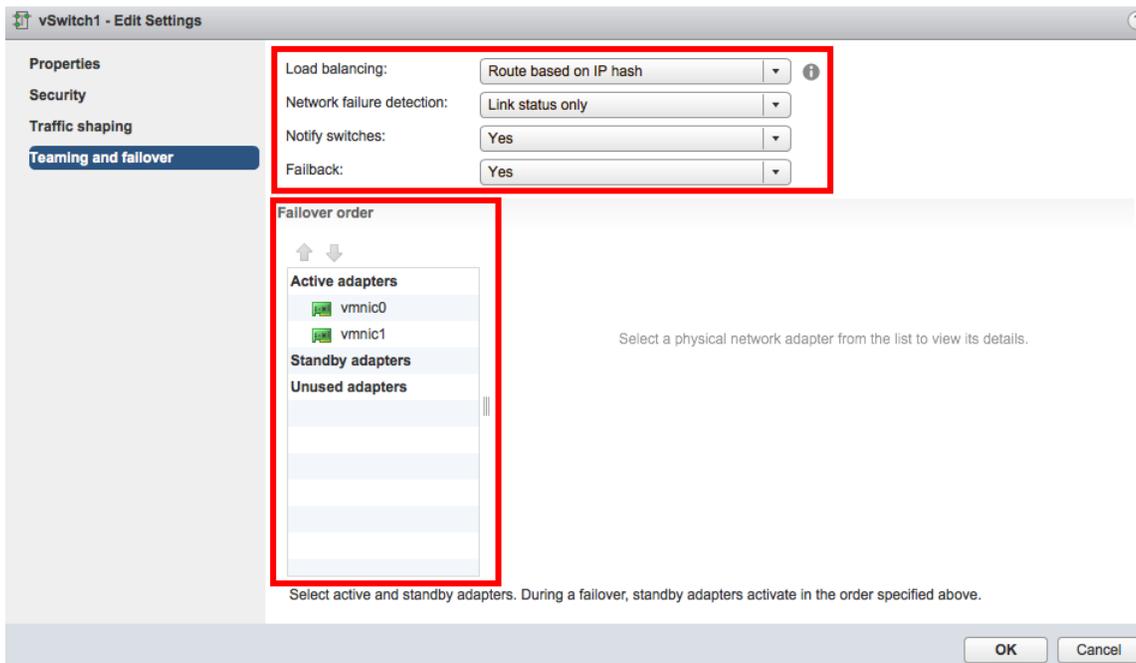


Figure 10. VMware vSphere virtual switch NIC teaming configuration shown on VMware vSphere VCenter Server Web Client

**Note:** Since VMware vSphere 5.1, VMware has added the LACP feature on vSphere hosts utilizing vSphere Distributed Switch (VDS). Use of VMware vSphere VDS in combination with VMware direct I/O technology and pass-through-capable hardware is recommended. Performance gains have been reported when using the combination of these technologies. However, the examples that follow do not use pass-through-capable hardware, and so configuration of these features is beyond the scope of this paper. For more information about this technology, refer to “Overview of vNetwork Distributed Switch Concepts” (VMware KB 1010555), which is listed in “Appendix B: References.”

Figure 11 through Figure 13 show three certified and recommended network topologies that provide fault tolerance and reliability between an Oracle ZFS Storage Appliance cluster configuration and VMware vSphere 6.x hosts.

**Note:** NFS direct-attached topology is not part of the VMware vSphere official certification program for storage vendors, so the NFS direct-attached topology for Oracle ZFS Storage Appliance systems is not officially certified.

Figure 11 shows the combination of port channel technology with the IEEE 802.3ad LACP topology between an Oracle ZFS Storage Appliance cluster configuration and a VMware vSphere 6.x host.

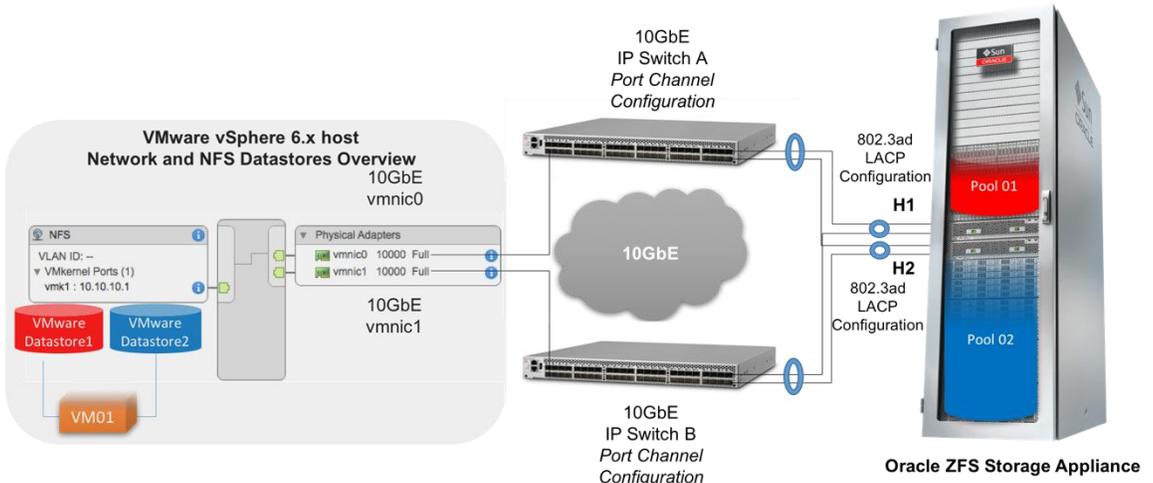


Figure 11. Combination of port channel technology with IEEE 802.3ad LACP network topology for an Oracle ZFS Storage Appliance cluster configuration and VMware vSphere 6.x host

Figure 12 shows the IPMP topology between an Oracle ZFS Storage Appliance cluster configuration and a VMware vSphere 6.x host.

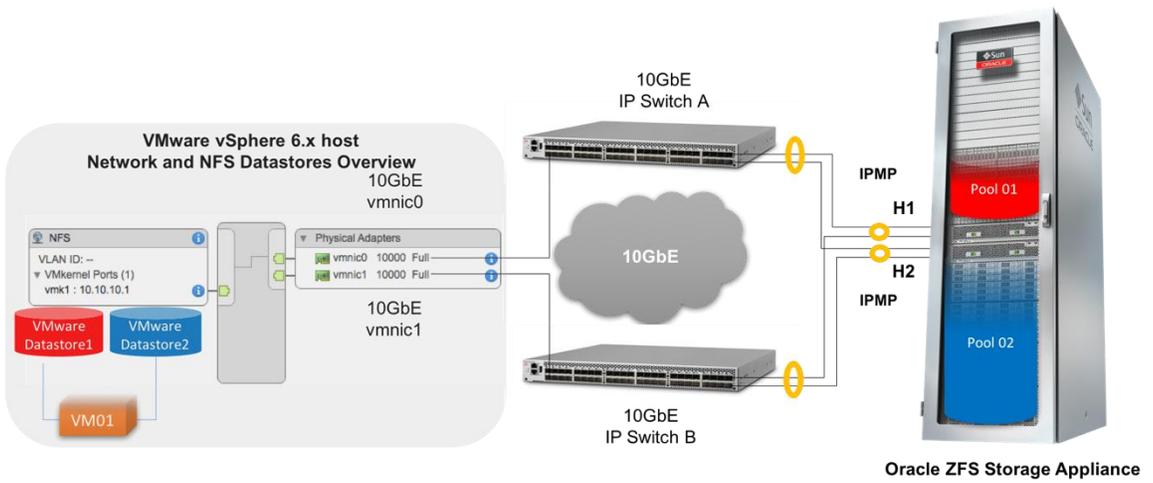


Figure 12. IPMP network topology between a VMware vSphere 6.x host and an Oracle ZFS Storage Appliance cluster configuration

Figure 13 shows another example of the IPMP topology between an Oracle ZFS Storage Appliance cluster configuration and a VMware vSphere 6.x host.

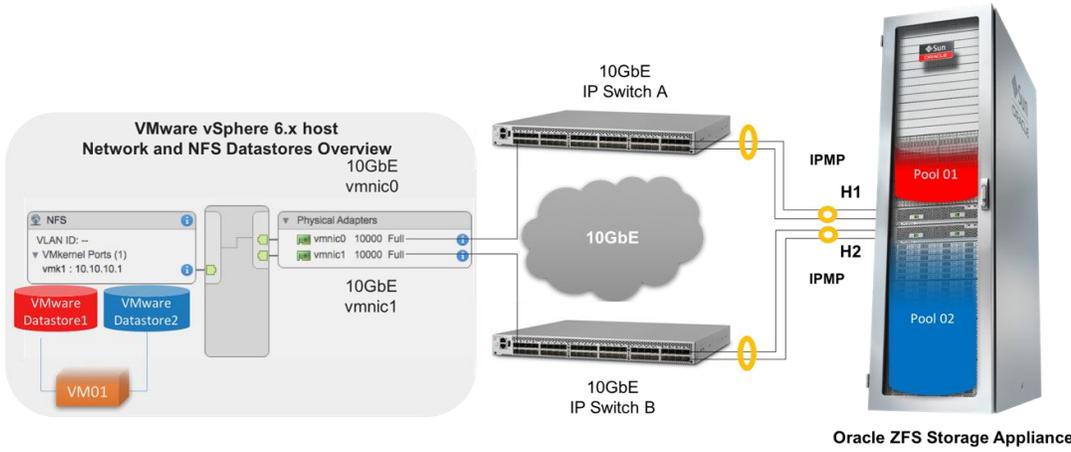


Figure 13. Combination of port channel technology with IEEE 802.3ad LACP and an Oracle ZFS Storage Appliance system and a VMware vSphere 6.x host

### Recommendations for the NFS Protocol

As a best practice, be sure to validate and alter the NFS and TCP/IP advanced settings prior to starting tests or putting the VMware vSphere 6.x hosts with the Oracle ZFS Storage Appliance cluster into production. The NFS and TCP/IP advanced settings are extremely important to ensure fault tolerance, resiliency, and high availability of the VMware NFS datastores during the Oracle ZFS Storage Appliance cluster failback and takeover operations.

Table 2 shows the VMware advanced settings that need to be altered. To alter the listed parameters, go to VMware vSphere Web Client 6.x and then select a VMware vSphere 6.x host. Select the **Manage** and **Settings** tabs and select **System > Advanced System Settings**. Use the search window to locate each of the advanced settings listed in Table 2, and then alter the values as recommended. See Figure 14.

**Note:** This configuration must be performed in all VMware vSphere 6.x host members of the VMware vSphere cluster. A reboot of each VMware vSphere 6.x host is needed to activate the new settings.

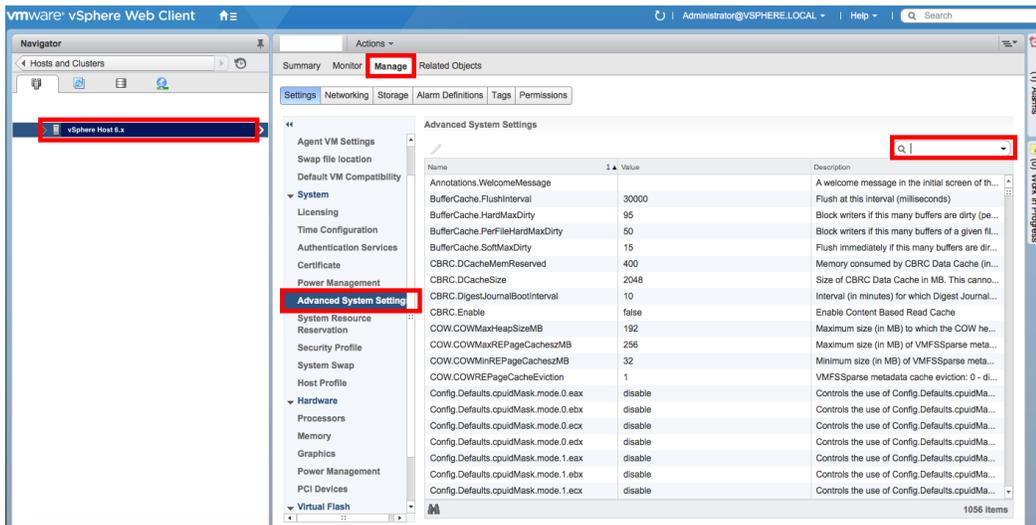


Figure 14. VMware vSphere advanced system settings shown in VMware vSphere Web Client

**TABLE 2: RECOMMENDED NFS AND TCP/IP ADVANCED SETTINGS FOR VMWARE VSPHERE 6.X DATASTORES WITH ORACLE ZFS STORAGE APPLIANCE CLUSTER CONFIGURATION**

VMware vSphere Advanced Settings Option	VMware vSphere Advanced Settings Value	VMware vSphere Advanced Settings Description
NFS.HeartbeatTimeout	5	Time in seconds before an outstanding heartbeat is stopped.
NFS.SendBufferSize	264	Default size of socket send buffer (KB)
NFS.ReceiveBufferSize	256	Default size of socket receive buffer (KB)
NFS.MaxVolumes	256	Maximum number of mounted NFS v3 volumes allowed per VMware vSphere 6.x host
Net.TcpipHeapMax	512	Max megabytes the TCP/IP module heap can grow to
Net.TcpipHeapSize	32	Size of the TCP/IP module heap in megabytes
NFS.HeartbeatFrequency	20	Time in seconds between heartbeats
NFS.HeartbeatDelta	12	Time in seconds since the last successful update before a heartbeat was sent
NFS.HeartbeatMaxFailures	10	Number of sequential failures before a volume is marked down

### VMware vSphere 6.x Cluster Recommendations

VMware vSphere 6.x cluster configuration is beyond of the scope of this white paper; however, when working with Oracle ZFS Storage Appliance systems, the following options are recommended:

- » Work with the VMware vSphere high availability (HA) and vSphere Distributed Resources Scheduled (DRS) cluster options.
- » At the cluster automation level, use the “fully automated” option and choose the priority level that best fits your virtualized environment.
- » For the power management cluster (DPM), choose the automatic option and select the DPM threshold that best fits your virtualized environment.
- » Enable host monitoring options and admission control.
- » Choose the virtual machine restart option for the cluster. The example reflects the “VM restart medium priority” and “Powered on for host isolation response” options.
- » Enable the VM monitoring option and choose the sensitivity level.
- » Enable the “Enhanced vMotion Compatibility” option for the VMware vSphere cluster. Choose the right VMware vSphere EVC mode for the hosts’ CPU (AMD or Intel).
- » For better high availability management of a VMware vSphere 6.x cluster, and also to avoid false positives due to network problems, enable the VMware vSphere Heartbeating feature for a VMware vSphere cluster with at least two nodes.

## Virtual Machine and Data Layout Recommendations

Recommendations for virtual machines and data layout as well as best practices for a VMware virtual machine working with Oracle ZFS Storage Appliance cluster configuration are as follows:

- » Work with VMware Virtual Machine version 11 or greater.
- » To improve network performance among virtual machines, datastores, and Oracle ZFS Storage Appliance systems, use the VMXNET3 network adapter.  
For a Linux virtual machine, select VMware Paravirtual (PVSCSI) as the default virtual SCSI controller for the virtual machine. PVSCSI provides lower CPU cost per I/O operation, and it is recommended for use cases and VMware Linux virtual machines and environments with intensive I/O workloads.
- » For virtual machines with the Microsoft Windows 2008 operating system release (or newer), use the LSI Logic SAS virtual SCSI controller.
- » For virtual machines or virtual desktops running operating systems older than Windows 2008, use LSI Logic Parallel virtual SCSI.
- » Install VMware client tools. For additional information regarding VMware tools installation or upgrade, refer to “Installing and Upgrading VMware Tools in vSphere” (VMware KB 2004754), which is listed in “Appendix B: References.”
- » When working with the Microsoft Windows platform, ensure that the latest service pack as well as recommended patches have been installed.
- » As a best practice and to achieve better performance for virtualized applications as well as management of the virtual environment, work with a multipool design with multiple datastores repositories in VMware vSphere 6.x.

Refer to the [VMware vSphere blog](#) for additional information regarding virtual SCSI recommendations.

Figure 15 shows the high-level view of a virtual machine layout with a multipool design.

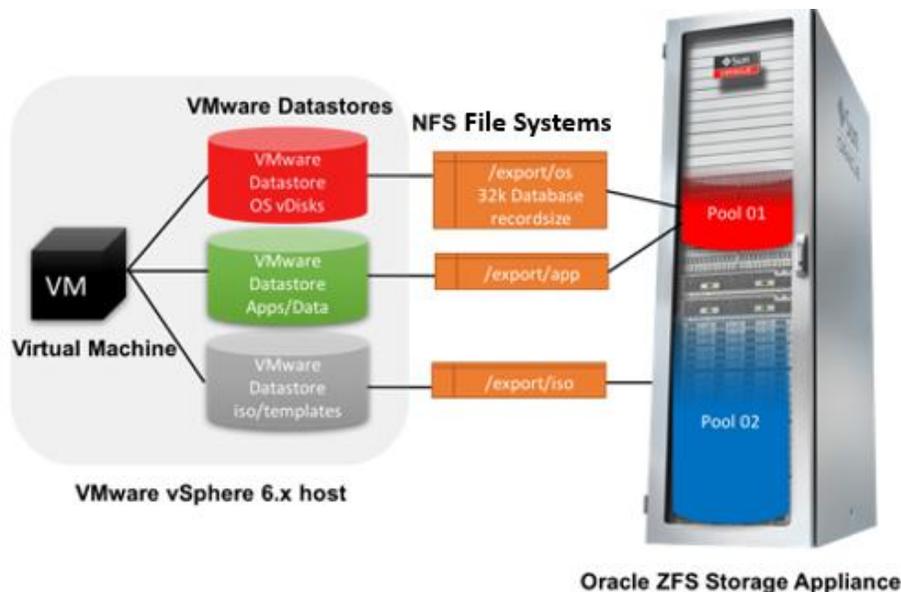


Figure 15. Recommended data layout for a VMware virtual machine with an Oracle ZFS Storage Appliance system

## Monitoring VMware vSphere 6.x with DTrace Analytics and the esxtop Tool

DTrace Analytics is a powerful and advanced monitoring tool provided by Oracle ZFS Storage Appliance systems. DTrace Analytics gives the storage administrator unique visibility into the entire system: monitoring in real time of different statistics for the operating system stack, storage resources, and protocols used as well as the I/O throughput and performance of the virtualized environment.

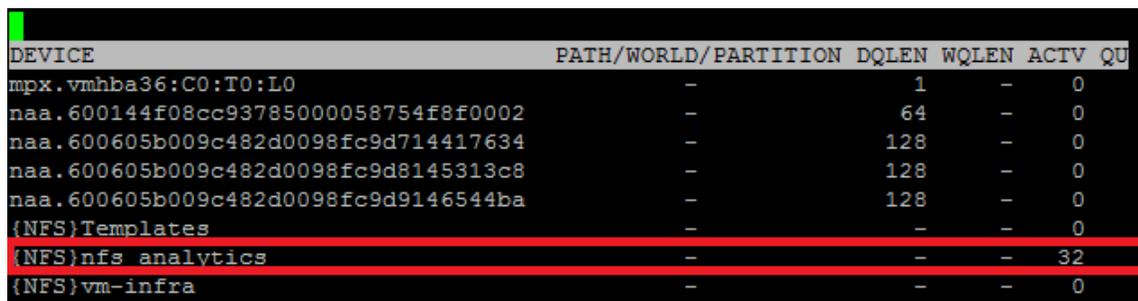
VMware provides a powerful monitoring tool called esxtop that is used at the VMware ESXi host level to monitor the performance and resource usage of the virtual environment. With this tool, you can identify possible bottlenecks, I/O performance issues, and network degradation as well as throughput levels.

For the most realistic report, VMware esxtop and DTrace Analytics should always be used together to validate as well as monitor your entire VMware storage performance and throughput. To ensure that your VMware NFS configuration is properly working, use the following DTrace Analytics and esxtop options.

### Monitoring NFS Performance

The following figures illustrate different examples of VMware esxtop and DTrace Analytics outputs used in monitoring NFS datastores' utilization and performance as well as the IP network. The NFSv4.1 protocol is used for virtual machines' disk datastores. DTrace Analytics is monitoring the virtual machines' disk usage in IOPS for each vmdk file.

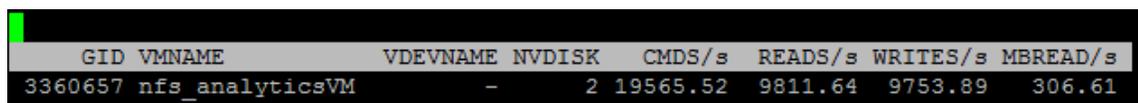
Figure 16 shows the VMware esxtop option for monitoring NFS datastores. To perform this, run esxtop, then type `u`. To alter the update time to every 2 seconds, type `s`, type `2` for seconds, and press Enter.



DEVICE	PATH/WORLD/PARTITION	DQLEN	WQLEN	ACTV	QU
mpx.vmhba36:C0:T0:L0	-	1	-	0	
naa.600144f08cc93785000058754f8f0002	-	64	-	0	
naa.600605b009c482d0098fc9d714417634	-	128	-	0	
naa.600605b009c482d0098fc9d8145313c8	-	128	-	0	
naa.600605b009c482d0098fc9d9146544ba	-	128	-	0	
{NFS}Templates	-	-	-	0	
{NFS}nfs_analytics	-	-	-	32	
{NFS}vm-infra	-	-	-	0	

Figure 16. Monitoring NFS datastores with VMware esxtop

Figure 17 shows the VMware esxtop option for monitoring virtual machines. To perform this, run esxtop, then type `v`. To alter the update time to every 2 seconds, type `s`, type `2` for seconds, and press Enter. The example shows the virtual machine `nfs_analyticsVM`.



GID	VMNAME	VDEVNAME	NVDISK	CMDS/s	READS/s	WRITES/s	MBREAD/s
3360657	nfs_analyticsVM	-	2	19565.52	9811.64	9753.89	306.61

Figure 17. Monitoring VMware virtual machine with VMware esxtop

The following figures show different examples of DTrace Analytics that are used in combination with VMware esxtop for monitoring network performance and throughput.

Figure 18 shows DTrace Analytics options for monitoring network datalinks as well as interface bytes and TCP bytes.



Figure 18. Monitoring network datalinks, interface bytes, and TCP bytes with DTrace Analytics

Figure 19 and Figure 20 show the DTrace Analytics options for monitoring the NFS protocol broken down by type of operation, clients, shares, and file name—in this case, .vmdk files.

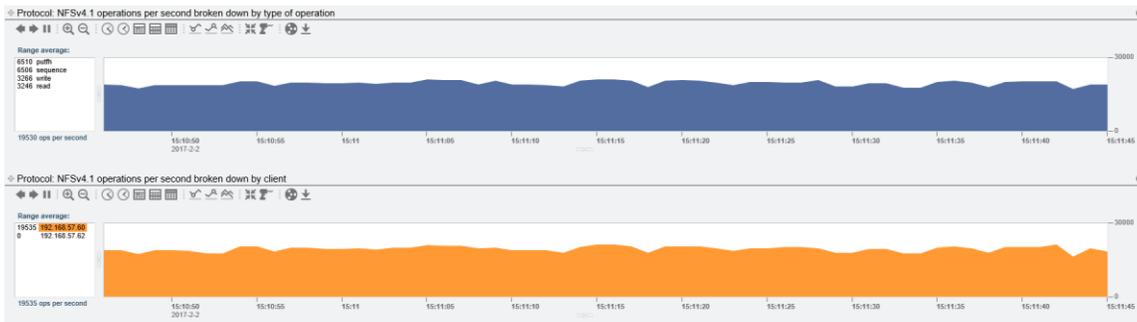


Figure 19. Monitoring the NFS protocol broken down by type of operation and client with DTrace Analytics

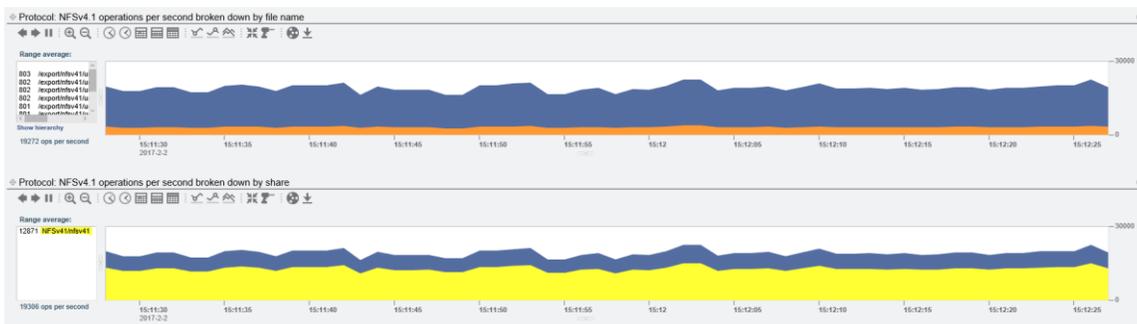


Figure 20. Monitoring the NFS protocol broken down by file name and share with DTrace Analytics

Figure 21 shows more DTrace Analytics options for monitoring the NFS protocol operations broken down by latency and by size, as well as cache ARC broken down by hit/miss.



Figure 21. Monitoring the NFS protocol operations broken down by latency and size, and cache ARC broken down by hit/miss with DTrace Analytics

## Conclusion

Oracle ZFS Storage Appliance systems offers outstanding performance for virtualized environments. The systems' architectural features and intelligent caching technology are designed to boot thousands of virtual machines, deliver thousands of IOPS for them, and provide the best throughput and response time for virtualized environments at a lower cost and as the simplest back-end storage solution (no additional tiering software or hardware is required).

VMware is a robust hypervisor that offers an easy way to manage your virtualized infrastructure. In combination, the VMware and Oracle ZFS Storage Appliance platforms and technologies are an excellent choice for your virtualized environments.

## Appendix A: Benchmark Results

Refer to the following websites for further information on testing results for Oracle ZFS Storage Appliance systems:

- » [SPC Benchmark 2 Executive Summary](#)
- » [SPC Benchmark 2 Full Disclosure Report](#)
- » [Active SPC-2 Results for Oracle ZFS Storage Appliance systems](#)

Also see the [Oracle quality awards for NAS](#).

## Appendix B: References

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

**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 Appliance products. Some cited documentation might still carry these legacy naming conventions.

- » Oracle ZFS Storage Appliance documentation library, including installation, analytics, customer service, and administration guides: [docs.oracle.com/en/storage/](https://docs.oracle.com/en/storage/)  
The *Oracle ZFS Storage Appliance Administration Guide* is also available through the Oracle ZFS Storage Appliance help function. The Help function in Oracle ZFS Storage Appliance systems can be accessed through the BUI.
- » Oracle ZFS Storage Appliance software releases and updates, available from My Oracle Support at [support.oracle.com](https://support.oracle.com) under Oracle Support Document 2021771.1
- » Oracle ZFS Storage Appliance Plug-in Downloads page: [oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html](https://oracle.com/technetwork/server-storage/sun-unified-storage/downloads/zfssa-plugins-1489830.html)
- » Oracle storage product information: [oracle.com/us/products/storage/overview/index.html](https://oracle.com/us/products/storage/overview/index.html)
- » Oracle ZFS Storage Appliance technical white papers and solution briefs on the Oracle Technology Network website, including "Best Practices for Oracle ZFS Storage Appliance and VMware vSphere 5.x": [oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html](https://oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html)
- » VMware: [vmware.com](https://vmware.com) and [vmware.com/support/pubs/](https://vmware.com/support/pubs/)
- » "Tuning ESXi for better storage performance by modifying the maximum I/O block size" (VMware KB 1003469): [kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=1003469](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1003469)
- » "Installing and upgrading VMware Tools in vSphere" (VMware KB 2004754): [kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=2004754](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2004754)
- » "vSphere Networking": [pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-networking-guide.pdf](https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-networking-guide.pdf)
- » "Overview of vNetwork Distributed Switch Concepts" (VMware KB 1010555): [kb.vmware.com/selfservice/microsites/search.do?language=en\\_US&cmd=displayKC&externalId=1010555](https://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1010555)



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

**Integrated Cloud Applications & Platform Services**

Copyright © 2017, 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. 0417

Best Practices for Configuring Oracle ZFS Storage Appliance and VMware vSphere 6.x with NFS Protocol  
April 2017  
Author: Anderson Souza, Application Integration Engineering  
Contributing author: Nithya Subramanian, Application Integration Engineering