ORACLE

# Best Practices of VMware vSphere 8.x Fibre Channel Protocol with Oracle ZFS Storage

Best practices and recommendations of VMware vSphere 8.x utilizing Fibre Channel protocol with Oracle ZFS Storage

Version [1.0]

Public

ORACLE

## Purpose statement

This document provides the best practices and recommendations of VMware vSphere 8.x utilizing Fibre Channel protocol with Oracle ZFS Storage.

## Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

ORACLE

# Table of contents

ORACLE

# Introduction

This white paper provides best practices and recommendations for configuring VMware vSphere 8.x with Oracle ZFS Storage to reach the optimal I/O performance and throughput. The outlined best practices and recommendations highlight configuration and tuning options for Fibre Channel for a VMware vSphere 8.x environment working with an Oracle ZFS Storage. In addition, this white paper presents the Fibre Channel architecture available with the Oracle ZFS Storage, how that architecture functions, and the overall setup process for utilizing FC LUNs on the system.

This content is provided for informational purposes and self-supported guidance only. Consultancy or other assistance related to the content is not covered under the Oracle Support contract or associated service requests. If you have questions or additional needs, then please do reach out to your Oracle Sales contact directly.

## Fibre Channel Architecture of the Oracle ZFS Storage

The supported Fibre Channel Host Bus Adapters (HBAs) for the Oracle ZFS Storage all have two ports. This means that each unit provides two connections to a Fibre Channel SAN infrastructure. Using a dual-fabric-based SAN design ensures full path redundancy to the host. The configuration described in this paper details an Oracle ZFS Storage cluster using a dual-fabric SAN infrastructure with VMware vSphere 8.x as shown in the following diagram.

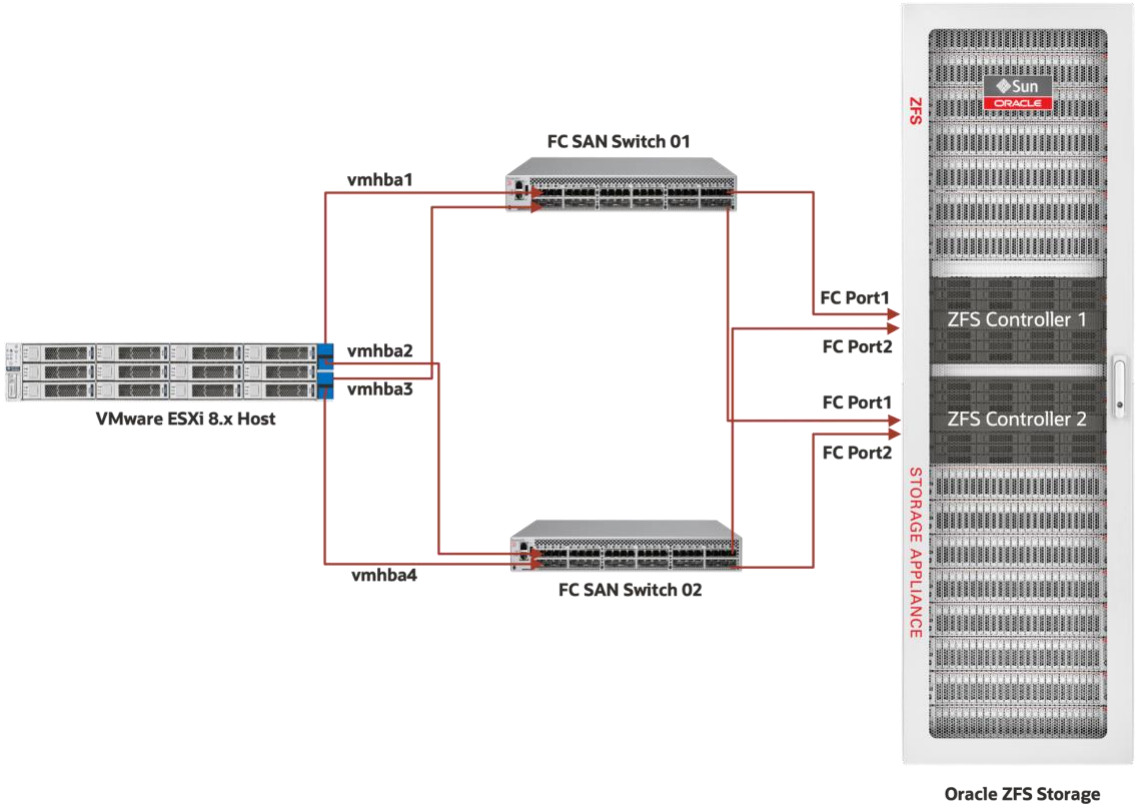


Figure 1. Fibre Channel connection diagram of VMware vSphere 8.x with Oracle ZFS Storage

To ensure no single point of failures in the SAN components/paths, both host HBA port connections always 'see' both Oracle ZFS Storage nodes. For each FC LUN on the Oracle ZFS Storage, each host HBA path sees both an active connection and a standby connection. Each Oracle ZFS Storage node runs an instance of an embedded Oracle Solaris operating system. The multipath element to handle I/O redirection from failed I/O paths to remaining active path is part of the Common Multi-protocol SCSI Target (COMSTAR) framework. All configuration information within the COMSTAR framework, like target and initiator group(s) information, is kept in sync between the two Oracle ZFS Storage nodes by the cluster functionality in the nodes. LUNs configured in a pool on the Oracle ZFS Storage appear on FC ports on both nodes of the Oracle ZFS Storage cluster. The two paths to the node that owns the pool are visible to the host as Active paths, while the path to the LUN through the node that does not own the pool has Standby status. This means a host can initiate a failover of the traffic between active paths to a LUN, like a dual FC connection to an Oracle ZFS Storage node. A failover between an Active and Standby path can be initiated by an Oracle ZFS Storage node failover. All pools must fail over to the node that had the Standby path status before the node failover.

## Storage Profile for VMware vSphere 8.x

Virtual machines infrastructures produce high random I/O patterns and need high storage performance as well as availability, low latency, and fast response time. To meet these demands, use a mirrored data profile on Oracle ZFS Storage. 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 SSDs' log devices and the Oracle ZFS Storage architecture, this profile can produce a large amount of input/output operations per second (IOPS) to attend to the demand of critical virtualized environments.

## Fibre Channel Setup and Configuration Sequence

To make the Fibre Channel installation and configuration process as easy as possible, execute the following steps in the listed sequence. More details for these steps are provided in following sections.

1. **Oracle ZFS Storage**: Set the Oracle ZFS Storage HBAs up for target mode and identify the World Wide Names (WWNs) of the FC ports. These are needed when setting up the zones in the SAN switch(es).
2. **Fibre Channel host connections:** Verify the FC HBA(s) on the host(s) are recognized by VMware vSphere 8.x (OS) level and identify the WWNs of the host HBA FC ports.
3. **SAN configuration and zoning setup:** Make the correct cable connections among the switch(es), ESXi hosts and Oracle ZFS Storage cluster. Verify all connection lights come up, indicating an established link with the switch at the desired speed.
4. **Configure zoning using HBA port and Oracle ZFS Storage FC HBA port WWNs.**
5. **Configure the FC target and initiator groups on Oracle ZFS Storage nodes.** To define access to the Oracle ZFS Storage nodes using the WWNs of each of the host HBA ports, target and initiator groups must be configured on the Oracle ZFS Storage nodes.

## Fibre Channel Zoning Configuration for VMware vSphere 8.x

Assuming that both the initiator and target FC-HBAs have been configured, proper routing in the SAN must be defined by means of SAN zones. The preferred zoning configuration often depends on a company's IT security policies and management rules. VMware vSphere 8.x with Oracle ZFS Storage Appliance supports different zoning methodology, for example: Single initiator (ESXi host side) to multiple targets (Oracle ZFS Storage Appliance side). Multiple initiators (ESXi host side) to single targets (Oracle ZFS Storage Appliance side), or single initiator (ESXi host side) to single target (Oracle ZFS Storage Appliance side), less utilized use case though.

In the following example, we used a single initiator to multiple targets zoning configuration. The ESXi 8.x host has two dual HBA fibre channel cards, and Oracle ZFS Storage Appliance has one dual Fibre Channel card per storage controller. Make sure you have at least two paths to each storage controller (Oracle ZFS Storage node). This redundancy avoids outages if a SAN connection between switch and storage controller fails – a situation that would not initiate an automatic node failover.

The following SAN configuration ensures optimal availability between VMware vSphere 8.x hosts and Oracle ZFS Storage Appliance

NOTE: Before creating zones, make sure the SAN configuration is correctly cabled.

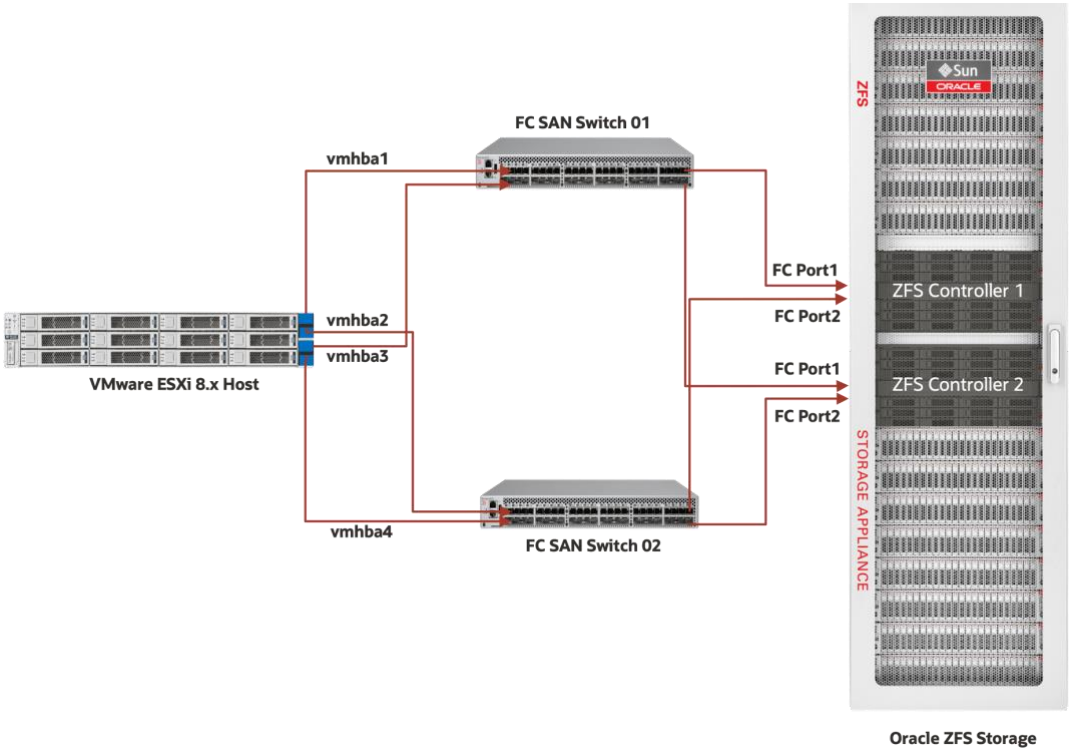


Figure 2. Fibre Channel zoning configuration recommended for VMware vSphere 8.x with Oracle ZFS Storage

| FC Zoning Configuration | | |
|---|---|---|
| VMware ESXi 8.x HBAs | Oracle ZFS Storage Controller 1 FC Ports | Oracle ZFS Storage Controller 2 FC Ports |
| Vmhba1 | FC Port 1 | FC Port 1 |
| Vmhba2 | FC Port 2 | FC Port 2 |
| Vmhba3 | FC Port 1 | FC Port 1 |
| Vmhba4 | FC Port 2 | FC Port 2 |

Table 1. Fibre Channel zoning configuration example

With the zones properly set up, you can now configure FC target and initiator groups on the Oracle ZFS Storage. The construct of target groups is used within the Oracle ZFS Storage's interface to allocate LUNs to certain FC ports of target FC HBAs in the Oracle ZFS Storage.

A target group acts as a pool of target FC ports from which LUNs are made visible to the outside world. Initiator groups in the Oracle ZFS Storage are used to control host (initiators) access to the LUNs. So, an initiator group acts as a pool of initiator FC ports that can be given access to a LUN in the Oracle ZFS Storage. Start by setting the Fibre Channel ports to be used in the Oracle ZFS Storage to Target mode, as seen in the pull-down menu selection in the following figure. To have an active failover path within the Oracle ZFS Storage node and a node service failover path to the other node, all four target ports of the Oracle ZFS Storage cluster need to be configured in the same target group as shown in the following screenshot.

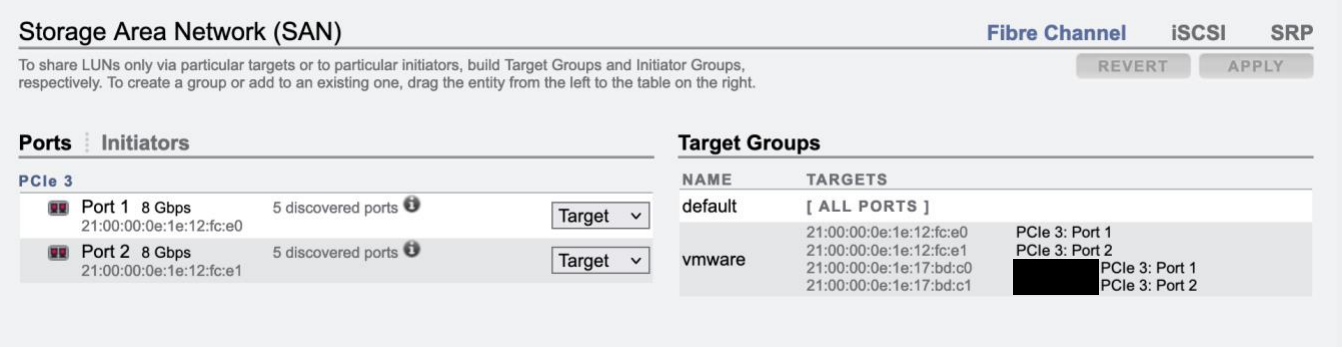


Figure 3. Oracle ZFS Storage SAN Target Groups configuration

Next, host ports are allocated in an initiator group. The initiator group is used later in the setup process, when LUNs are created. Provide logical, meaningful names for the initiator ports.

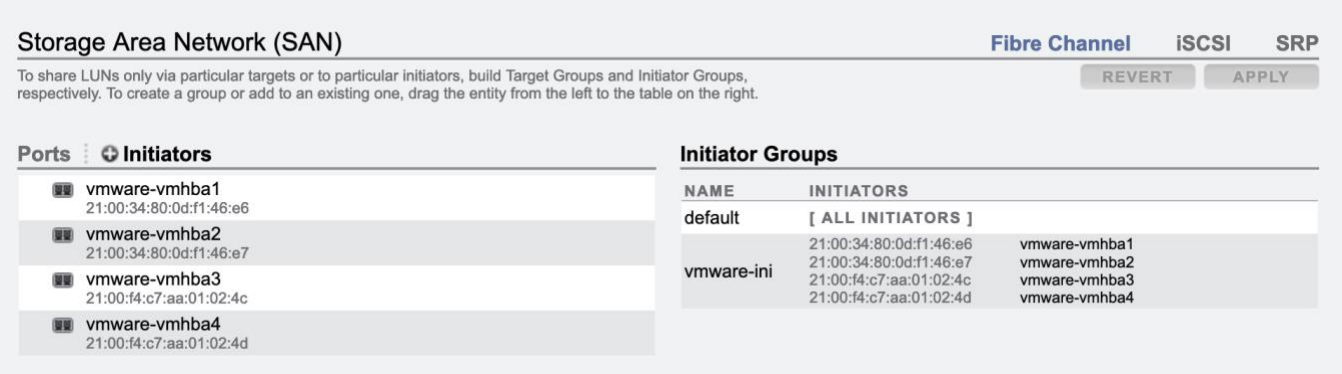


Figure 4. Oracle ZFS Storage SAN Initiators Groups configuration

## LUN Creation and Configuration for VMware vSphere 8.x

In this session, we will list the best practices and recommendations for Fibre Channel LUNs on the Oracle ZFS Storage for VMware vSphere 8.x environments.

- Ensure that the host bus adapter (HBA) is on the VMware HCL
- Ensure you have the latest Oracle ZFS Storage Appliance software release. Refer to Oracle ZFS Storage Appliance: Software Updates (Doc ID 2021771:1)
- Update the Fibre Channel host bus adapters' (HBAs') firmware and drivers to their latest version
- Ensure that your Storage Area Network (SAN) has been designed for high availability and load balance without critical points of failure
- Work with at least two Fibre Channel switches and at least one dual port 32Gbps HBA per Oracle ZFS Storage controller and VMware ESXi8.x host
- Verify the visibility of configured LUNs on the host(s)
- Detect the GUID (Global unique identifier) in the device name of the LUN by matching it with the GUID presented by the Oracle ZFS Storage BUI
- Verify the SAN zoning configuration
- Create LUNs and set up the membership of target and initiator groups to the LUNs
- Verify that the LUN attributes, like block size and cache behavior requirements on the Oracle ZFS Storage node, are set as required. Attributes can be changed at any time, however, once you have created a LUN, its block size attribute

cannot be changed. Set the LUN block accordingly for the type of your workload. Listed below is the list of recommended block sizes:

| Workload Characterization | Recommended Block Size |
|---|---|
| General purpose workload | 64k |
| Vswap | 64k |
| Microsoft Exchange DB | 32k |
| Microsoft Logs | 128k |
| Linux Operating System | 64k |
| 100% OLTP DB | 8K |
| 100% OLAP Workload | 128k |
| Windows Boot disks | 64k |

- Enable Thin provisioned for all LUNs utilized by VMware vSphere 8.x
- Enable Space Reclamation feature on the LUN. This option will allow the utilization of SCSI Space reclamation feature. The command helps thin-provisioned storage arrays to reclaim unused space from the VMFS datastore and thin virtual disks on the datastore. Refer to Space Reclamation on vSphere VMFS Datastores for more information on how to setup space reclamation on VMFS6.
- Leave read and write limit as default
- Leave data deduplication feature disabled
- Enable compression. Select LZJB compression algorithm
- Set checksum to Fletcher4 (Standard)
- Set additional replication to Normal (Single Copy)
- For cache device usage, select All data and metadata option
- Set synchronous write bias to latency
- Enable prevent share destruction

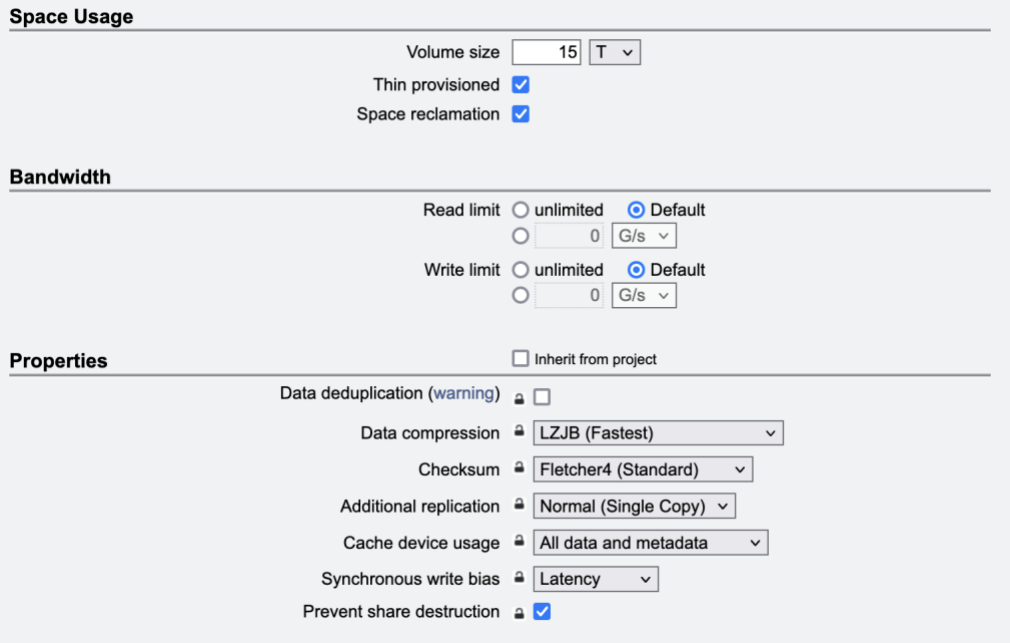

Figure 5. Fibre Channel LUN configuration recommended for VMware vSphere 8.x on Oracle ZFS Storage

On VMware vSphere 8.x host side:

- Ensure you have only one VMware VMFS (virtual Machine File System) volume per LUN
- For raw devices, use RDM (Raw Device Mapping)

## VMware Path Selection Policies (PSP) and Storage Array Type Plugin (SATP)

On VMware ESXi, Path Selection Policies (PSPs) determine how I/O traffic is distributed across multiple available paths between the ESXi host and a storage device. PSPs are part of the VMware Native Multipathing (NMP) system and work in conjunction with Storage Array Type Plugins (SATPs) to provide redundancy, load balancing, and failover in storage environments.

Storage Array Type Plugin (SATP) is a component of the VMware Native Multipathing (NMP) system. SATPs are responsible for managing specific behaviors and failover policies of different types of storage arrays. They ensure ESXi can effectively handle multipathing for various storage systems, allowing for redundancy, load balancing, and failover capabilities in SAN and iSCSI environments.

Each SATP works in conjunction with Path Selection Policies (PSPs) to manage how the ESXi host interacts with storage devices over multiple paths. The SATP ensures that the storage array's specific characteristics (like Active/Active, Active/Passive, or Asymmetric Logical Unit Access - ALUA) are properly handled.

The VMware Path Selection Policies (PSPs) recommended for Oracle ZFS Storage is VMW_PSP_RR (Round Robin). This policy alternates I/O requests across all available paths, balances the load across paths to optimize performance and it is typically used in Active/Active storage array configurations and is beneficial for spreading the workload.

The VMware Storage Array Type Plugin (SATP) recommended for Oracle ZFS Storage is VMW_SATP_ALUA (Asymmetric Logical Unit Access). This will allow Oracle ZFS Storage to expose which paths are optimized (preferred) or non-optimized.

## Setting up Path Selection Policies (PSP) and Storage Array Type Plugin (SATP) on VMware vSphere 8.x

Assuming you already have all Fibre Channel LUNs presented to your VMware ESX8.x host, follow the step by step below to properly setup the correct Path Selection Policies (PSP) and Storage Array Type Plugin (SATP) for Oracle ZFS Storage Fibre Channel Luns

1. Identify the correct LUNs from Oracle ZFS Storage with the following esxcli command line listed below:

**NOTE:** The example shows the following command line for capturing only Oracle ZFS Fibre Channel disks as well as changing the path selection policy.

**esxcli storage nmp device list | egrep -i "SUN Fibre Channel Disk"**

```
Device Display Name: SUN Fibre Channel Disk (naa.600144f0fccbbec8000066e0a3890001)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa94f000b)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa8ff0009)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aab40000d)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa8d70008)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa8930006)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa8b50007)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aaa77000c)
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa92a000a)
```

2. Check and adjust the storage array type of the Oracle ZFS Storage LUNs to VMW_SATP_ALUA and the path selection policy to VMW_PSP_RR

**esxcli storage nmp satp set --default-psp=VMW_PSP_RR --satp=VMW_SATP_ALUA**
**esxcli storage nmp device list**

```
naa.600144f0fccbbec8000066e0a3890001
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa92a000a)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: {implicit_support=on;explicit_support=off;
explicit_allow=on;alua_followover=on;{TPG_id=0,TPG_state=AO}}
Path Selection Policy: VMW_PSP_MRU
Path Selection Policy Device Config: Current Path=vmhba7:C0:T0:L6
Path Selection Policy Device Custom Config:
Working Paths: vmhba7:C0:T0:L6
```

**esxcli storage nmp device list | egrep -i "SUN Fibre Channel Disk" | awk '{ print $8 }' | cut -c 2-37**

```
naa.600144f0fccbbec8000066e0a3890001
naa.600144f0c36f708b0000509aa94f000b
naa.600144f0c36f708b0000509aa8ff0009
naa.600144f0c36f708b0000509aab40000d
naa.600144f0c36f708b0000509aa8d70008
naa.600144f0c36f708b0000509aa8930006
naa.600144f0c36f708b0000509aa8b50007
naa.600144f0c36f708b0000509aaa77000c
naa.600144f0c36f708b0000509aa92a000a
```

3. Ensure that you are not using round robin path selection before performing changes.

**for a in `esxcli storage nmp device list | egrep -i "SUN Fibre Channel Disk" | awk '{ print $8 }' | cut -c 2-37`**
**do**
**esxcli storage nmp psp roundrobin deviceconfig get -d $a**
**done**

```
Device naa.600144f0fccbbec8000066e0a3890001 Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aa94f000b Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aa8ff0009 Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aab40000d Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aa8d70008 Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aa8930006 Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aa8b50007 Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aaa77000c Does not use the Round Robin path selection policy.
Device naa.600144f0c36f708b0000509aa92a000a Does not use the Round Robin path selection policy.
```

4. Run the following command to change the path selection policy VMW_PSP_MRU to VMW_PSP_RR:

**for a in `esxcli storage nmp device list | egrep -i "SUN Fibre Channel Disk" | awk '{ print $8 }' | cut -c 2-37`**
**do**

**esxcli storage nmp device set -d $a --psp=VMW_PSP_RR**
**done**

5. Run the following command to ensure that the new path selection policy has been updated:

**esxcli storage nmp device list**

```
naa.600144f0fccbbec8000066e0a3890001
Device Display Name: SUN Fibre Channel Disk (naa.600144f0c36f708b0000509aa92a000a)
Storage Array Type: VMW_SATP_ALUA
Storage Array Type Device Config: {implicit_support=on;explicit_support=off;
explicit_allow=on;alua_followover=on;{TPG_id=0,TPG_state=AO}}
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0;lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
Path Selection Policy Device Custom Config:
Working Paths: vmhba6:C0:T0:L6, vmhba7:C0:T0:L6
```

6. Change the I/O operation limit value to 1 and the type of the round robin path switching to iops for all Fibre Channel disks on the Oracle ZFS Storage. List the device configuration before changing.

**esxcli storage nmp psp roundrobin deviceconfig get -d naa.600144f0fccbbec8000066e0a3890001**

```
Byte Limit: 10485760
Device: naa.600144f0fccbbec8000066e0a3890001
IOOperation Limit: 1000
Limit Type: Default
Use Active Unoptimized Paths: false
```

Perform the configuration:

**for a in `esxcli storage nmp device list | egrep -i "SUN Fibre Channel Disk" | awk '{ print $8 }' | cut -c 2-37`**
**do**
**esxcli storage nmp psp roundrobin deviceconfig set -d $a -I 1 -t iops**
**done**

Check if the IO operation limit and limit type has been changed accordingly:

**esxcli storage nmp psp roundrobin deviceconfig get -d naa.600144f0fccbbec8000066e0a3890001**

Byte Limit: 10485760
Device: naa.600144f0fccbbec8000066e0a3890001
IOOperation Limit: 1
Latency Evaluation Interval: 0 milliseconds
Limit Type: Iops
Number Of Sampling IOs Per Path: 0
Use Active Unoptimized Paths: false

7. Run the following command to ensure that the new values for operation limit and round robin path switching have been updated:

```
for a in `esxcli storage nmp device list | egrep -i "SUN Fibre Channel Disk" | awk '{ print $8 }' | cut -c 2-37`
do
esxcli storage nmp psp roundrobin deviceconfig get -d $a
done
```

Device: naa.600144f0c36f708b0000509aa92a000a

8. As a best practice for VMware vSphere 8.x and Oracle ZFS Storage, adjust the queue depth option for all HBAs attached with the system. Use the following steps to accomplish this task. Identify which HBA module is currently loaded on the VMware hypervisor using the following commands.

For QLogic HBAs, run:

```
esxcli system module list | grep qla*
```
qlnativefc true true

Note: The example uses QLogic HBAs (module qlnativefc).

For Emulex HBAs, run:

```
esxcli system module list | grep lpfc*
```

9. Use the following commands to set a new queue depth value.

For QLogic HBAs, run:

```
esxcli system module parameters set -p ql2xmaxqdepth=64 -m qlnativefc
esxcli system module parameters list -m qlnativefc
```

The following is the output for QLogic HBAs:

Name Type Value Description

--------------------------- ---- ----- -----------------------------------------

ql2xmaxqdepth int 64 Maximum queue depth to report for target

10. For Emulex HBAs run the following esxcli command listed below. For this example, we are using an Emulex lpfc820 HBA model.

```
# esxcli system module parameters set -p lpfc0_lun_queue_depth=64 -m lpfc820
```

11. On VMware ESXi host web client, create a new VMware datastore on the Oracle ZFS Storage Fibre Channel LUNs. Click Storage, devices, new datastore, then enter a new datastore name. Click next.

**NOTE:** If you are working with VMware vSphere client, the steps to configure a new datastore might be different. Refer to VMware vSphere official documentation for details.

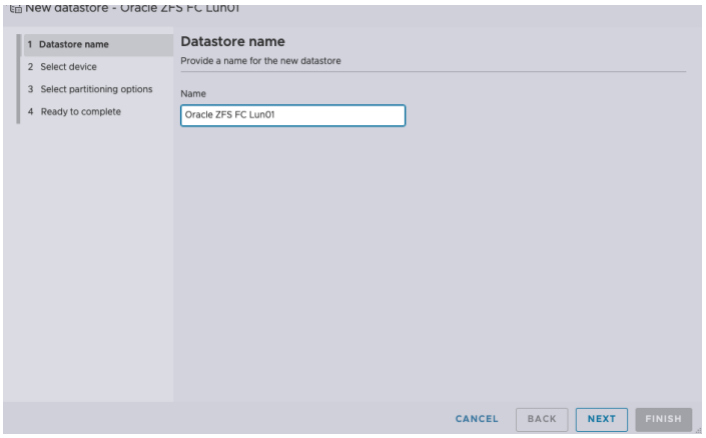


Figure 6. VMware vSphere 8.x datastore configuration using Oracle ZFS Storage Fibre Channel LUNs

12. Select the Oracle ZFS Storage Fibre Channel LUN to create a new VMFS partition.

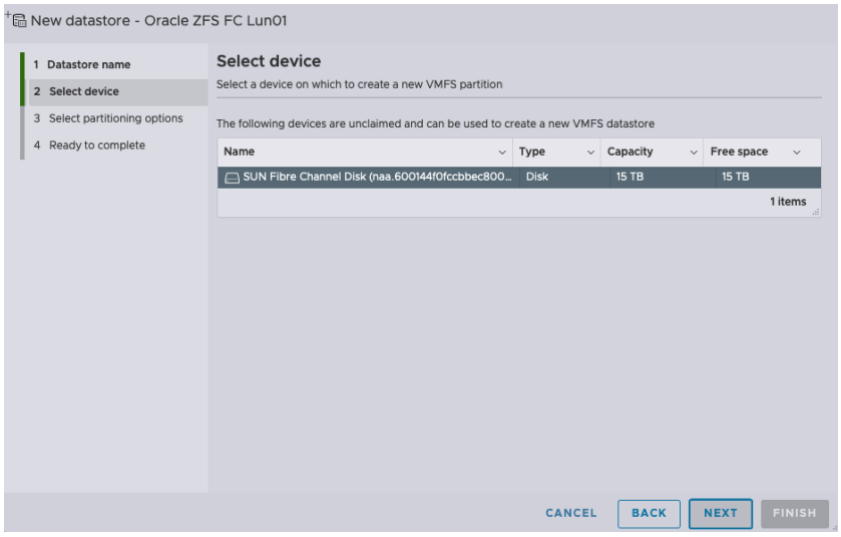


Figure 7. Oracle ZFS Storage Fibre Channel LUNs overview on VMware vSphere 8.x host.

13. Select partition options to best fit your environment. Choose VMFS version 6. Click next and finish to complete the VMware datastore configuration using Oracle ZFS Storage LUNs.

NOTE: For this example, we are using a full 15TB fibre channel LUN from Oracle ZFS Storage formatted as VMFS version 6.

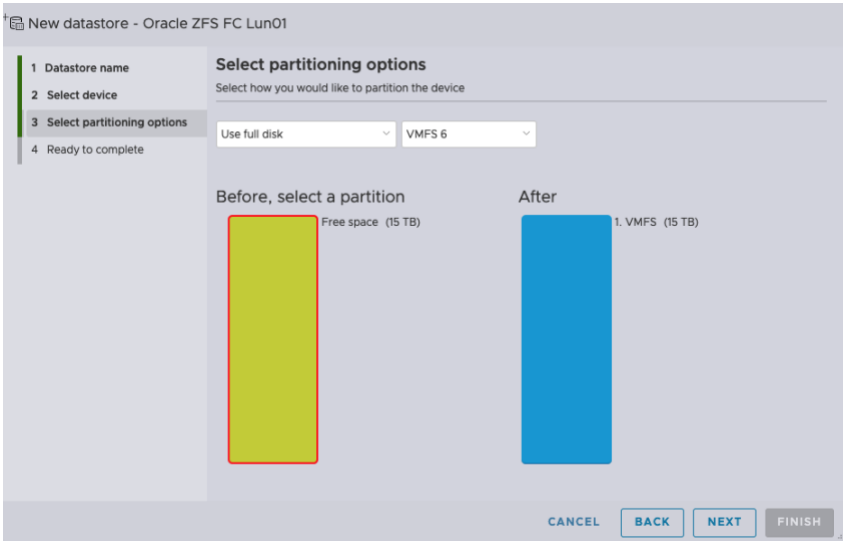

Figure 8. VMware vSphere 8.x VMFS partitioning

14. The new VMware datastore using Oracle ZFS Fibre Channel LUN will be created and available for virtual machines workloads.

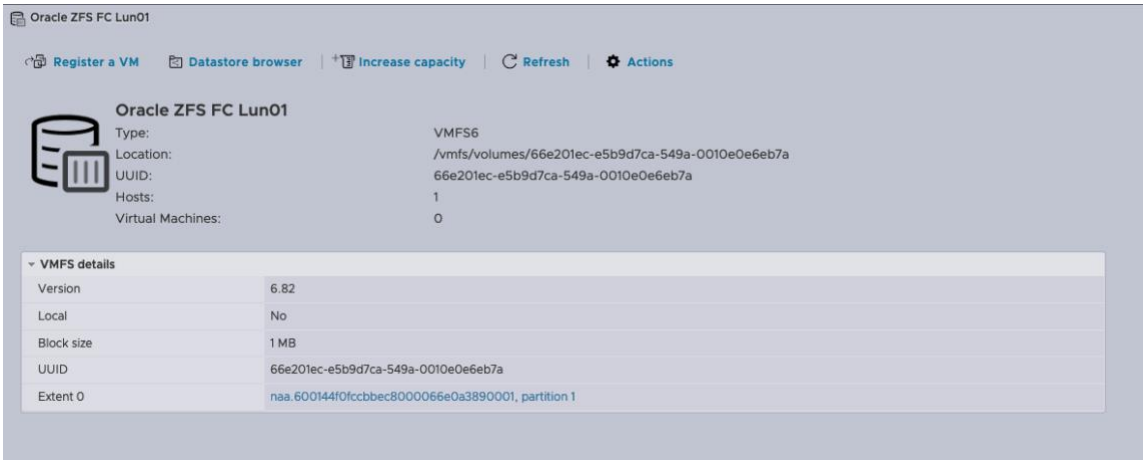


Figure 8. VMware vSphere 8.x datastore using Oracle ZFS Storage Fibre Channel LUNs

## Monitoring VMware with DTrace Analytics and ESXTOP

DTrace Analytics is a powerful and advanced monitoring tool provided by the ZFS Storage Appliance. 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 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 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.

NOTE: VMware ESXTOP and DTrace Analytics should always be used together to validate as well as monitor your entire VMware storage performance and throughput for the most realistic report.

## Monitoring Fibre Channel Performance

The following examples show how to use ESXTOP and DTrace Analytics to monitor VMware Fibre Channel LUNs, datastore, and HBA performance and throughput. For VMware ESXTOP, open a SSH connection with your ESXi8.x host and run the following commands:

1. Type esxtop and then press **d** for monitoring VMware HBAs. Option **d** or disk view option (HBA mode) can be used for monitoring virtual HBAs.

2. Press **s 2** to alter the update time to every **2** seconds, and press **Enter**.

NOTE: Ensure that the virtual HBAs (vmhbas) are correctly balancing the I/O. The example below highlights **vmhba1, vmhba2, vmhba3,** and **vmhba4** correctly load balancing the I/O.

| ADAPTR | PATH | NPTH | CMDS/s | READS/s | WRITES/s | MBREAD/s | MBWRTN/s | DAVG/cmd | KAVG/cmd | GAVG/cmd | QAVG/cmd |
|---|---|---|---|---|---|---|---|---|---|---|---|
| vmhba0 | - | 0 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| vmhba1 | - | 8 | 11060.16 | 5492.33 | 5567.83 | 343.24 | 347.94 | 1.42 | 1.76 | 3.18 | 0.00 |
| vmhba2 | - | 8 | 11057.22 | 5513.30 | 5543.92 | 344.50 | 346.34 | 1.42 | 1.75 | 3.17 | 0.00 |
| vmhba3 | - | 8 | 11249.35 | 5602.23 | 5647.12 | 350.13 | 352.84 | 1.40 | 1.76 | 3.16 | 0.00 |
| vmhba32 | - | 1 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| vmhba4 | - | 8 | 11268.22 | 5603.91 | 5664.31 | 350.24 | 353.96 | 1.39 | 1.76 | 3.16 | 0.00 |
| vmhba5 | - | 1 | 4.19 | 0.00 | 4.19 | 0.00 | 0.02 | 0.03 | 0.01 | 0.04 | 0.00 |
| vmhba64 | - | 0 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| vmhba65 | - | 0 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| vmhba66 | - | 0 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| vmhba67 | - | 0 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |

Figure 9. VMware ESXTOP output of the “d” option.

For additional esxtop disk monitoring performance metrics, type **f**. Type **s 2** to alter the update time to every **2** seconds and then press **Enter**.

```
* A:  ADAPTR = Adapter Name
* B:  PATH = Path Name
* C:  NPATHS = Num Paths
  D:  QSTATS = Queue Stats
* E:  IOSTATS = I/O Stats
  F:  RESVSTATS = Reserve Stats
* G:  LATSTATS/cmd = Overall Latency Stats (ms)
  H:  LATSTATS/rd = Read Latency Stats (ms)
  I:  LATSTATS/wr = Write Latency Stats (ms)
  J:  ERRSTATS/s = Error Stats
  K:  PAESTATS/s = PAE Stats
  L:  SPLTSTATS/s = SPLIT Stats

Toggle fields with a-l, any other key to return:
```

Figure 10. VMware ESXTOP monitoring performance metrics.

The following figures show different examples of DTrace Analytics that can be used in combination with VMware ESXTOP for monitoring Fibre Channel performance and throughput. The first DTrace Analytics metric for Fibre Channel is the Fibre Channel operations per second broken down by initiator. With this metric you can see the total IOPs from all VMware HBAs (vmhbas). The second important metric is the Fibre Channel operations per second broken down by LUN. This metrics can show us the total IOPs per second from all VMware LUNs. And the third metric, is the Fibre Channel operations per second broken down by type of operations. This metric will show your workload characterization broken down by the total of IOPs from reads and writes operations.

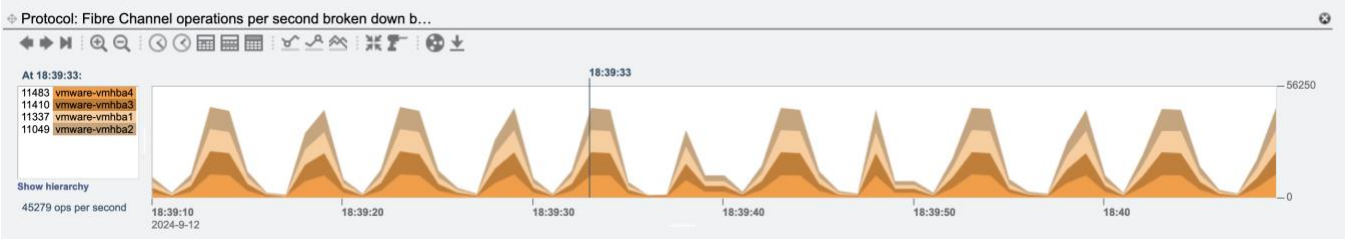


Figure 11. Fibre Channel operations per second broken down by initiator.

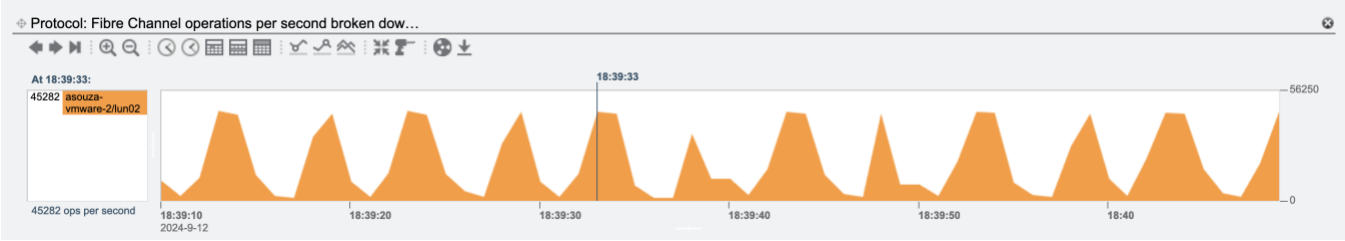


Figure 12. Fibre Channel operations per second broken down by LUN

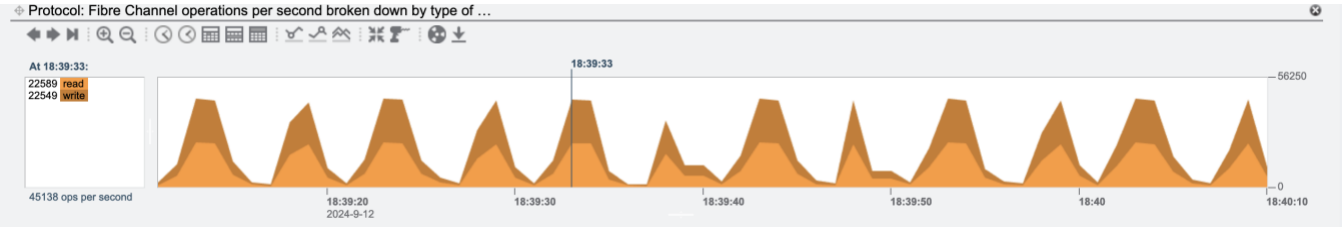


Figure 13. Fibre Channel operations per second broken down by type of operations

ORACLE

## Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle



**Author:** Anderson Souza