



ZFS STORAGE
APPLIANCE

An Oracle Technical White Paper
October 2013

Configuring VMware vSphere 5.1 with Oracle® ZFS Storage Appliance and Oracle Fabric Interconnect

An IP over InfiniBand configuration overview for VMware vSphere 5.1 with Oracle ZFS Storage Appliance and Oracle Fabric Interconnect

Introduction	2
Overview of Configuration Example's System Components	3
Oracle ZFS Storage Appliance Settings	4
Controllers, Software Release and Disk Pools	4
InfiniBand Network Settings for the Oracle ZFS Storage Appliance	6
Installing Oracle Virtual Networking Host Driver for VMware	10
Enabling IP over InfiniBand (IPoIB) Interfaces	11
Recommendations for VMware ESXi5.1 Multipathing Settings for iSCSI	16
Conclusion	18
Appendix A: Updating Mellanox HCA Firmware.....	19
Appendix B: References.....	23
Oracle ZFS Storage Appliance Documentation	23

Introduction

This white paper details how to configure a VMware ESXi 5.1 hypervisor with Oracle ZFS Storage Appliance and Oracle Fabric Interconnect using IP over InfiniBand protocol (IPoIB). The paper also provides the multipathing best practices and recommendation for configuring iSCSI and Fibre Channel LUNs in an IPoIB environment.

Highlighted in this paper are:

- VMware ESXi 5.1 with Oracle ZFS Storage Appliance and Oracle Fabric Interconnect environment overview
- How to download, install and enable the Oracle Virtual Networking host driver for VMware ESXi 5.1
- How to enable IP over InfiniBand VMware uplink network interfaces and configure VMware virtual switches
- Best practices and recommendations for iSCSI and Fibre Channel (FC) multipathing for IPoIB
- How to update the firmware of Mellanox InfiniBand host channel adapter (HCA) cards in a VMware ESXi 5.1 environment

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.

Overview of Configuration Example's System Components

The following tables describe the hardware configuration, operating systems, and software releases used in this white paper's configuration example.

Table 1 shows the hardware used.

TABLE 1. HARDWARE USED IN REFERENCE ARCHITECTURE

EQUIPMENT	QUANTITY	CONFIGURATION
Storage	1 cluster (2 controllers)	Oracle ZFS Storage ZS3-2 256 GB DRAM per controller 1 dual quad data rate (QDR) InfiniBand HCA M2 card per controller 4 x DE2-24C disk enclosure 3TB 7200RPM 2 x DE2-24P disk enclosure 900GB 10k RPM 1 dual port 10GbE network interface card (NIC) 1 dual port 8Gbps FC host bus adapter (HBA) 4 x 73GB Log device
Network	1	Oracle Fabric Interconnect model VP780-CH-QDR
Server	2	Oracle's Sun Blade X6275 M2 10GbE Server Module* 98GB DRAM 2 internal HDDs 1 x NEM switch 1Gbps 1 x dual HCA Mellanox Technologies MT26428 [ConnectX VPI - 10GigE / IB QDR, PCIe 2.0 5GT/s] – Firmware 2.11.2010

* Note that this server module is being updated to Oracle's Sun Blade X3-2B server module (formerly called the Sun Blade X6270 M3) and the new server module is fully compatible with the listed reference architecture.

Table 2 shows the virtual machine components used.

TABLE 2. VIRTUAL MACHINE COMPONENTS USED IN REFERENCE ARCHITECTURE

OPERATING SYSTEM	QUANTITY	CONFIGURATION
Oracle Linux 6.3	100	Standard Virtual Machine Guest

Table 3 shows the software used.

TABLE 3. SOFTWARE USED IN REFERENCE ARCHITECTURE

SOFTWARE	VERSION
Oracle® ZFS Storage Software	OS8
Oracle Virtual Networking Fabric Switch O.S	3.9.0
Vdbench	5.0.3
SWAT	305 Beta 21
VMware vCenter Server	5.1.0 (Build 880146)
VMware ESXi hypervisor software	5.1.0 (Build 799733)

Oracle ZFS Storage Appliance Settings

The following Oracle ZFS Storage Appliance configurations were used with VMware vSphere5.x and an Oracle Fabric Interconnect IPoIB environment.

Controllers, Software Release and Disk Pools

- Oracle ZFS Storage Appliance ZS3-2 cluster Active/Active mode
- Oracle ZFS Storage OS8 Software
- 2x8 cores 2.10GHz Intel® Xeon® CPU E5-2658 0 @ 2.10GHz per Oracle ZFS Storage Appliance controller
- 256GB of DRAM (ARC) L1 cache per Oracle ZFS Storage Appliance controller
- 1 x dual QDR InfiniBand HCA M2 card (4 physical standalone InfiniBand connections with Oracle Fabric Interconnect).
- A mirrored disk pool (capacity) composed of 84 disks for data and parity, 4 spare disks and 2 SSDs for log devices. Total of 111TB of usable disk space. See figure 2 for disk layout of capacity disk pool.
- A mirrored disk pool (performance) composed of 38 disks for data and parity, 2 spare disks and 2 SSDs for log devices. Total of 15TB of usable disk space. See figure 3 for disk layout of performance disk pool.

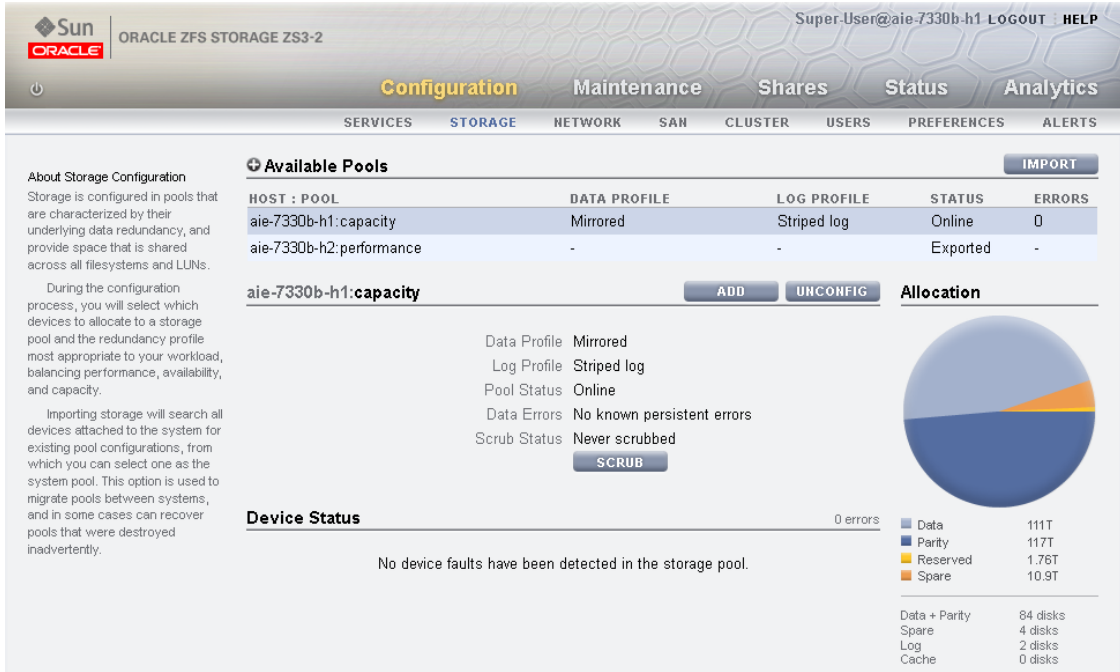


Figure 1. Capacity disk pool layout shown on Oracle ZFS Storage Appliance BUI

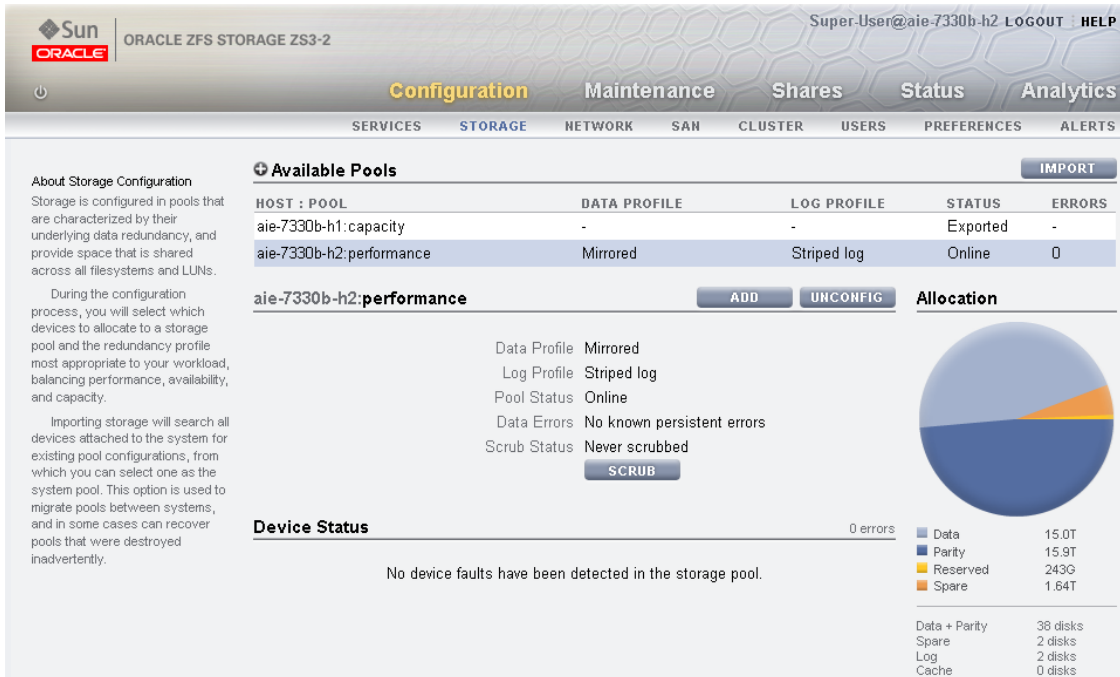


Figure 2. Performance disk pool layout shown on Oracle ZFS Storage Appliance BUI

Note: Additional information on an Oracle ZFS Storage Appliance cluster configuration can be found in the *Sun ZFS Storage 7000 System Administration Guide* at:

http://docs.oracle.com/cd/E26765_01/html/E26397/index.html

InfiniBand Network Settings for the Oracle ZFS Storage Appliance

This paper's example uses IP over InfiniBand (IPoIB) and standalone InfiniBand interfaces between Oracle ZFS Storage Appliance and VMware ESXi5.1 hosts. In the example, Oracle ZFS Storage Appliance has four physical InfiniBand interfaces connected to the Oracle Fabric Interconnect using InfiniBand connected mode link type (also known as CM link mode). For better performance and throughput, each InfiniBand interface has been configured with a full 40Gbps and four different IP addresses which will be presented to VMware ESXi5.1 hosts as Oracle ZFS Storage Appliance iSCSI targets. Also, Oracle Fabric Interconnect has four physical InfiniBand connections with VMware ESXi5.1 hosts which have Mellanox ConnectX-2 HCA InfiniBand cards.

Only the standalone InfiniBand configuration is supported for this environment, so an IPMP configuration for the Oracle ZFS Storage Appliance is not supported.

Oracle Fabric Interconnect provides a host driver for VMware ESXi5.1 that will be responsible for enabling the InfiniBand and IPoIB modules in the VMware kernel, so you will be able to configure 40Gbps uplink interfaces to work with Oracle ZFS Storage Appliance. The release 5.3.1 host driver for VMware ESXi5.1 has been qualified and certified with Oracle ZFS Storage Appliance to work with FC, iSCSI and NFS protocols.

Note: There is no Oracle Virtual Networking host driver for Oracle ZFS Storage Appliance. For this environment, a host driver for Oracle ZFS Storage Appliance is not needed. In this example scenario, the connection among Oracle Fabric Interconnect, VMware ESXi 5.1 hosts and the appliance will be the same as a standard InfiniBand network topology.

Oracle Fabric Interconnect provides configuration options for virtual network cards (vNICs) and virtual Fibre Channel adapters (vHBAs). These vNICs and vHBAs act exactly like any other NIC and HBA and can be dynamically provisioned in the Oracle Fabric Interconnect CLI environment or using Oracle Fabric Manager.

Because Oracle Fabric Interconnect provides a built-in subnet manager that is part of the firmware and is responsible for assigning local identifiers (LIDs) to each port connect to the InfiniBand fabric, as well as the configuration of the routing table based off the assigned LIDs, you do not need a different or standalone subnet manager installation for managing your InfiniBand network.

The following figure shows the high-level architecture of the IP over InfiniBand environment for VMware ESXi5.1 with Oracle Fabric Interconnect and Oracle ZFS Storage Appliance. In this architecture you will have options to work with an end-to-end IPoIB solution between VMware ESX.5.1 hosts and Oracle ZFS Storage Appliance.

Note: The configuration of virtual network interfaces (vNICs) and virtual host bus adapters (vHBAs) is beyond the scope of this white paper. For more detailed information on these configuration elements, please refer to the Oracle Virtual Interconnect documentation at:

http://docs.oracle.com/cd/E38500_01/

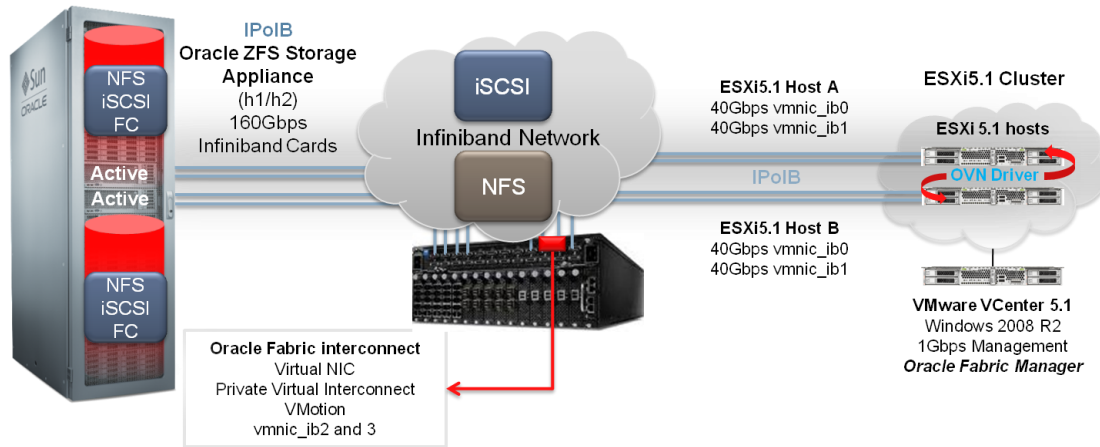


Figure 3. High-level architecture of VMware ESXi5.1 IPoIB environment with Oracle Fabric Interconnect and Oracle ZFS Storage Appliance

On the Oracle ZFS Storage Appliance side, there is one dual InfiniBand card per controller, configured as shown in figures 4, 5, 6 and 7.

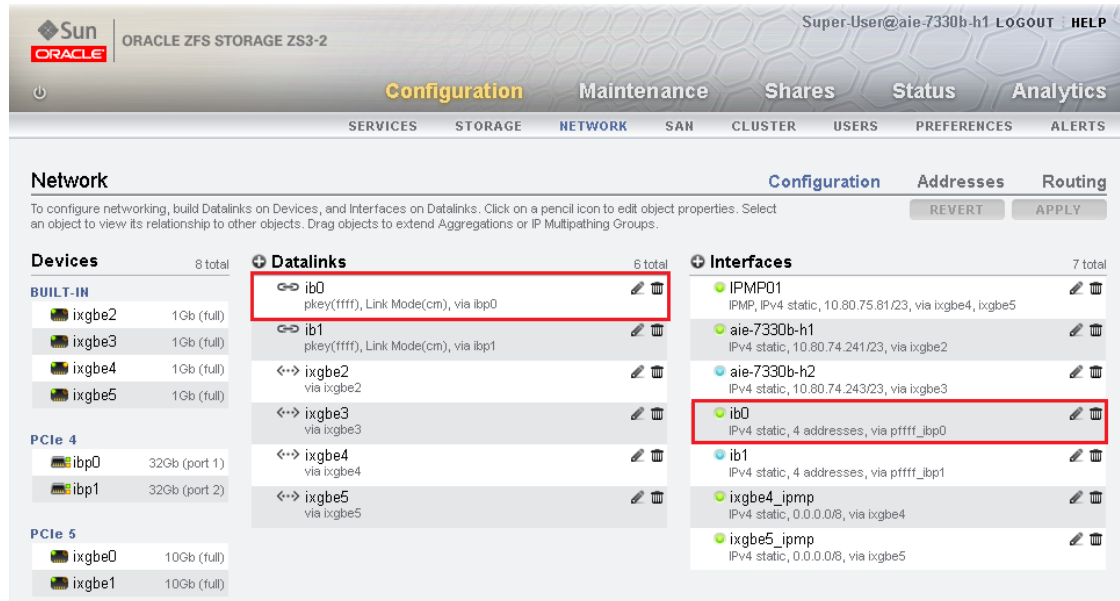


Figure 4. InfiniBand interface ib0 network configuration shown on Oracle ZFS Storage Appliance BUI

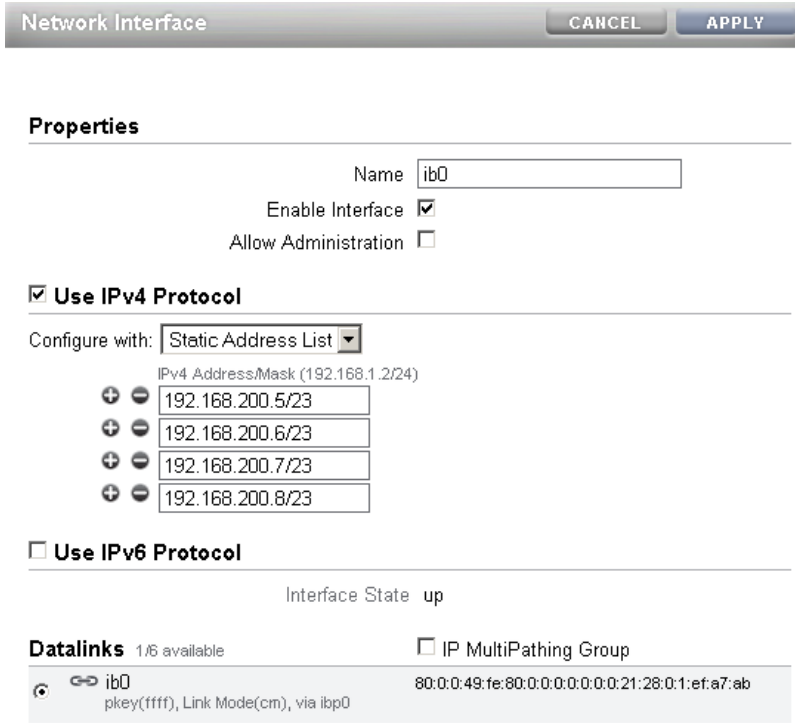


Figure 5. InfiniBand interface ib0 IP addresses configuration shown on Oracle ZFS Storage Appliance BUI

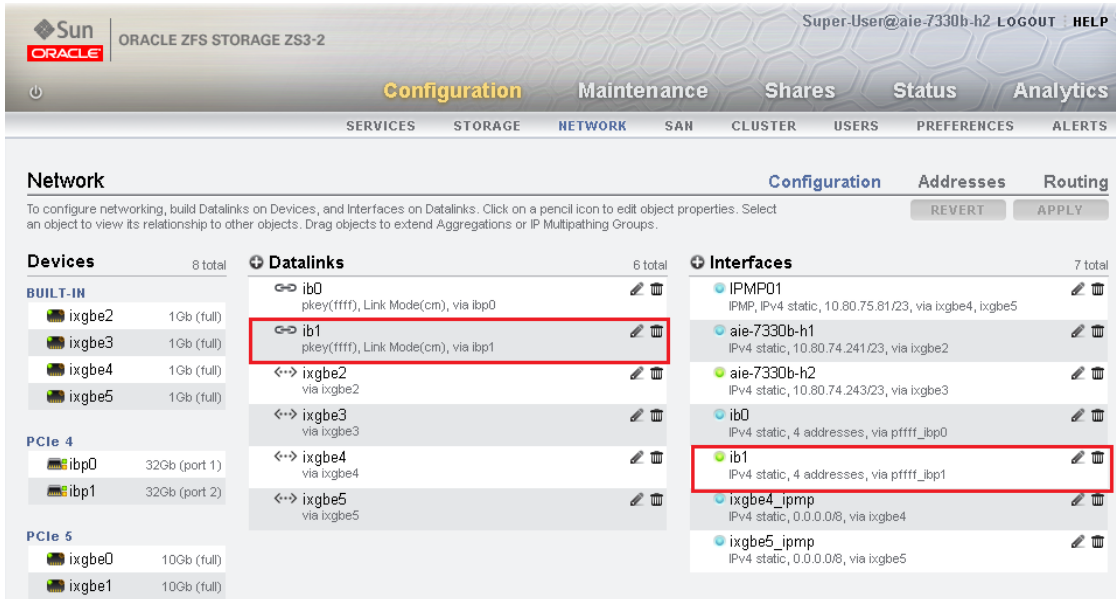


Figure 6. InfiniBand interface ib1 network configuration shown on Oracle ZFS Storage Appliance BUI

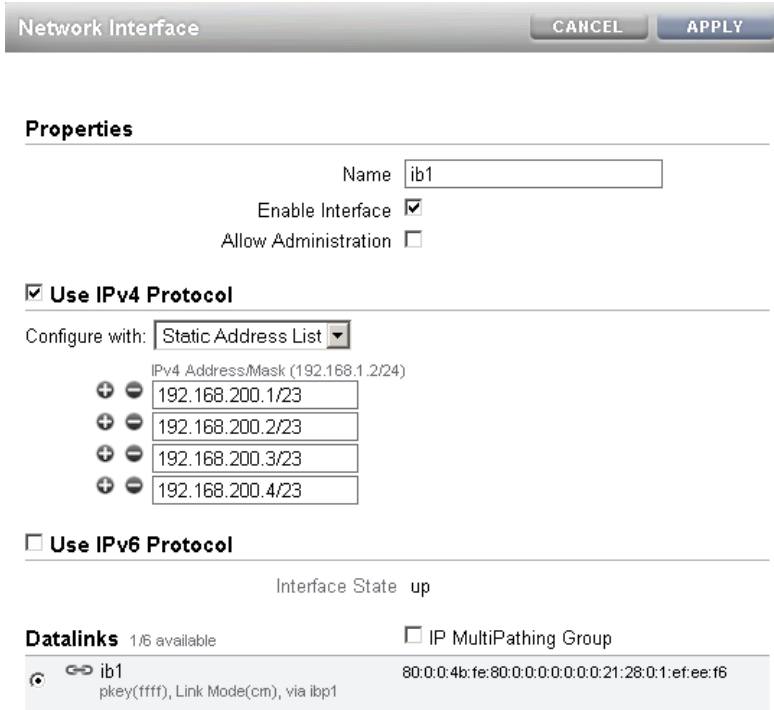


Figure 7. InfiniBand interface ib1 IP addresses configuration shown on Oracle ZFS Storage Appliance BUI

As part of the solution, Oracle Fabric Interconnect provides a host driver for VMware ESXi5.1 hosts that will be responsible for enabling the InfiniBand uplink physical interfaces in the VMware kernel as well as the IPoIB communication with Oracle Fabric Interconnect.

The following example shows the output from executing the `show physical-server` command line in Oracle Fabric Interconnect. In the example, the hosts `aie-6275m2a0` and `aie-6275m2a1` are working with the ESXi-5.3.1 host driver and the HCA firmware 2.11.2010.

Note: Steps for upgrading the HCA firmware and installing the Oracle Virtual Networking driver in a VMware environment are listed in the next sections of this white paper.

```
admin@qdr1[xsigo] show physical-server aie-6275m2a0
name      guid          descr          port          os              version          server-profile
-----
aie-6275m2a0  21280001efe456  qdr1:ServerPort12  VMware/ESXi-5.3.1.ESX.1/x86_64  2.11.2010/3.0.0
aie-6275m2a0  21280001efe455  qdr1:ServerPort13  VMware/ESXi-5.3.1.ESX.1/x86_64  2.11.2010/3.0.0
2 records displayed
admin@qdr1[xsigo]
```

Figure 8. Output of `show physical-server` command line shown on Oracle Fabric Interconnect CLI

Refer to the following link for additional, detailed information about the Oracle Fabric Interconnect ESX host driver as part of the full documentation set for the Oracle Fabric Interconnect.

http://docs.oracle.com/cd/E38500_01/

IMPORTANT: Oracle ZFS Storage Appliance IPMP (IP Multipathing Group) configuration is not supported in this environment. Instead, use the InfiniBand standalone physical network configuration between Oracle ZFS Storage Appliance and Oracle Fabric Interconnect.

Installing the Oracle Virtual Networking Host Driver for VMware

Follow these steps to install the Oracle Virtual Networking host driver in a VMware ESXi 5.1 host.

1. Download the Oracle Virtual Networking host driver for VMware ESXi5.1 from the following link:

<http://www.oracle.com/technetwork/server-storage/xsigo-1870185.html>

Xsigo Downloads, GPL

The following software is licensed pursuant to the General Public License version 2 (the "GPL"). The GPL may be viewed at this [location](#) or in the downloaded software.

Host Drivers

↓ Oracle Virtual Networking Drivers for Oracle Enterprise Linux
(md5sum: 2260a26ccce0036f1b5e9e625b75a4ab)

↓ Oracle Virtual Networking Drivers for RedHat Enterprise Linux
(md5sum: 5f0b8195f013a835af04cdb28965a653)

↓ Oracle Virtual Networking Drivers for VMware ESX 5.1
(md5sum: 7d9c6d0a83aa6e3bf15b08e1077ade02)

↓ Oracle Virtual Networking Drivers for VMware ESX 5
(md5sum: c4a4c789488f20f6866d5c2f2d11b638)

↓ Oracle Virtual Networking Drivers for VMware ESX 4.1
(md5sum: 38d1ddd83b9b420dc7dbccbe1203c0a)

↓ Oracle Virtual Networking Drivers for Oracle VM
(md5sum: 415bbd633b6af657a8b7755fbec5755f)

↓ OVM 3.2.1 Drivers version 5.0.6
(md5sum: fc70310c6ef5a159c88378b1ea2d2682) [Release Notes](#)

HCA FW Utils

↓ Oracle Virtual Networking Tool to burn option ROM 4.0.3 for Linux
(md5sum: 0a954e494712a00db0b128674184705e)

↓ Oracle Virtual Networking Tool to burn option ROM 4.0.2 for Microsoft Windows
(md5sum: 37ec1517ea9034bead92a667eca45a06)

Figure 9. Download link for Oracle Virtual Networking host driver for VMware ESXi5.1

2. Copy the .zip file to your ESXi5.1 hosts and install:

```
/tmp # esxcli software vib install -d /tmp/xsigo_5.3.1.ESX.1-1vmw.510.0.0.613838.zip
```

Note: Ensure that you have included the full path to the .zip file. The example uses the /tmp folder.

The screen output indicates that the installation completed successfully, but the system needs to be rebooted for the changes to be effective.

```

Reboot Required: true
VIBs Installed: Mellanox_bootbank_net-ib-core_5.3.1.ESX.1-
1OEM.510.0.0.613838, Mellanox_bootbank_net-ib-ipoib_5.3.1.ESX.1-
1OEM.510.0.0.613838, Mellanox_bootbank_net-ib-mad_5.3.1.ESX.1-
1OEM.510.0.0.613838, Mellanox_bootbank_net-ib-sa_5.3.1.ESX.1-
1OEM.510.0.0.613838, Mellanox_bootbank_net-mlx4-core_5.3.1.ESX.1-
1OEM.510.0.0.613838, Mellanox_bootbank_net-mlx4-ib_5.3.1.ESX.1-
1OEM.510.0.0.613838, VMware_bootbank_net-xscore_5.3.1.ESX.1-
1vmw.510.0.0.613838, VMware_bootbank_net-xsvnic_5.3.1.ESX.1-
1vmw.510.0.0.613838, VMware_bootbank_net-xve_5.3.1.ESX.1-1vmw.510.0.0.613838,
VMware_bootbank_scsi-xsvhba_5.3.1.ESX.1-1vmw.510.0.0.613838
VIBs Removed:
VIBs Skipped:

```

3. Reboot the ESXi5.1 hosts and execute the following command line to ensure that the Oracle Virtual Networking host driver for VMware has been successfully installed.

```
/tmp # esxcli software vib list | grep -i 5.3.1.ESX.1
```

```

/tmp # esxcli software vib list | egrep -i "PartnerSupported|Mellanox"
net-ib-core          5.3.1.ESX.1-1OEM.510.0.0.613838 Mellanox PartnerSupported 2013-07-02
net-ib-ipoib         5.3.1.ESX.1-1OEM.510.0.0.613838 Mellanox PartnerSupported 2013-07-02
net-ib-mad           5.3.1.ESX.1-1OEM.510.0.0.613838 Mellanox PartnerSupported 2013-07-02
net-ib-sa            5.3.1.ESX.1-1OEM.510.0.0.613838 Mellanox PartnerSupported 2013-07-02
net-mlx4-core        5.3.1.ESX.1-1OEM.510.0.0.613838 Mellanox PartnerSupported 2013-07-02
net-mlx4-ib          5.3.1.ESX.1-1OEM.510.0.0.613838 Mellanox PartnerSupported 2013-07-02
net-xscore           5.3.1.ESX.1-1vmw.510.0.0.613838 VMware PartnerSupported 2013-07-02
net-xsvnic           5.3.1.ESX.1-1vmw.510.0.0.613838 VMware PartnerSupported 2013-07-02
net-xve              5.3.1.ESX.1-1vmw.510.0.0.613838 VMware PartnerSupported 2013-07-02
scsi-xsvhba          5.3.1.ESX.1-1vmw.510.0.0.613838 VMware PartnerSupported 2013-07-02
/tmp #

```

Figure 10. Output from esxcli software vib list showing on VMware CLI environment

Enabling IP over InfiniBand (IPoIB) Interfaces

After installing the VMware host driver for ESX5.1 hosts, you must enable the InfiniBand IPoIB uplink interfaces on the VMware ESXi5.1 hosts as part of the overall installation process. The following simple steps must be completed before any additional configuration can take place.

1. Log in to the ESXi5.x host as user root.
2. Run the `/sbin/xg_ctrl_ipoib enable` command.
3. Reboot the ESXi5.x hosts.

Note: The ESXi5.x hosts must be rebooted to complete the configuration.

To disable the IPoIB driver, use the `disable` command. Then reboot the ESXi5.x hosts.

```
/sbin/xg_ctrl_ipoib disable
reboot
```

On the ESXi5.x hosts CLI, execute the `esxcfg-nics -l` command to ensure that the new InfiniBand vmnics have been successfully created. the following figure shows the two new vmnics, called `vmnic_ib0` and `vmnic_ib1`.

```
~ # esxcfg-nics -l
Name      PCI      Driver      Link Speed Duplex MAC Address  MTU  Description
vmnic0    0000:07:00.00 e1000e      Down 0Mbps   Half 00:1e:68:2f:24:9a 1500 Intel Corporation 80003E52LAN Gigabit Ethernet Controller
vmnic1    0000:07:00.01 e1000e      Up    1000Mbps Full 00:1e:68:2f:24:9b 1500 Intel Corporation 80003E52LAN Gigabit Ethernet Controller
vmnic2    0000:0d:00.00 e1000e      Up    1000Mbps Full 00:c0:dd:0a:b0:2e 1500 Intel Corporation 82571EB Gigabit Ethernet Controller
vmnic3    0000:0d:00.01 e1000e      Up    1000Mbps Full 00:c0:dd:0a:b0:2f 1500 Intel Corporation 82571EB Gigabit Ethernet Controller
vmnic_ib0 0000:03:00.00 ib_ipoib    Up    40000Mbps Full 00:21:28:a1:47:2d 1500 Mellanox Technologies MT26428 [ConnectX VPI - 10GigE / IB QDR, PCIe 2.0 5GT/s]
vmnic_ib1 0000:03:00.00 ib_ipoib    Up    40000Mbps Full 00:21:28:a1:47:2e 1500 Mellanox Technologies MT26428 [ConnectX VPI - 10GigE / IB QDR, PCIe 2.0 5GT/s]
```

Figure 11. Output of the VMware `esxcfg-nics -l` command in the VMware CLI

Network Adapters						
Device	Speed	Configured	Switch	MAC Address	Observed IP ranges	
Intel Corporation 80003E52LAN Gigabit Ethernet Controller						
vmnic1	1000 Full	Negotiate	vSwitch0	00:1e:68:2f:24:9b	10.80.74.1-10.80.74.127	
vmnic0	Down	Negotiate	None	00:1e:68:2f:24:9a	None	
Intel Corporation 82571EB Gigabit Ethernet Controller						
vmnic3	1000 Full	Negotiate	None	00:c0:dd:0a:b0:2f	10.80.74.1-10.80.74.127	
vmnic2	1000 Full	Negotiate	None	00:c0:dd:0a:b0:2e	10.80.74.1-10.80.74.127	
Mellanox Technologies MT26428 [ConnectX VPI - 10GigE / IB QDR, PCIe 2.0 5GT/s]						
vmnic_ib1	40000 F...	Negotiate	vSwitch2	00:21:28:a1:47:2e	None	
vmnic_ib0	40000 F...	Negotiate	vSwitch2	00:21:28:a1:47:2d	None	

Figure 12. InfiniBand uplink interfaces shown on VMware vSphere client network adapters page

At this point you will be able to create a new VMware virtual switch (vSwitch) with the `vmnic_ib0` and `vmnic_ib1` InfiniBand over IP uplink interfaces as well as present Oracle ZFS Storage Appliance iSCSI LUNs and NFS shares to your VMware ESXi5.1 hosts. The following example shows the created vSwitch01 with a VMkernel port called `IBoIP`. To configure a vSwitch, use the following steps and graphic example.

1. In the VMware vSphere client, select the ESXi5.1 host for which you want to configure a new virtual switch. Select the **Configuration** tab, then **Networking** and **Add Networking** as seen in the following figure.

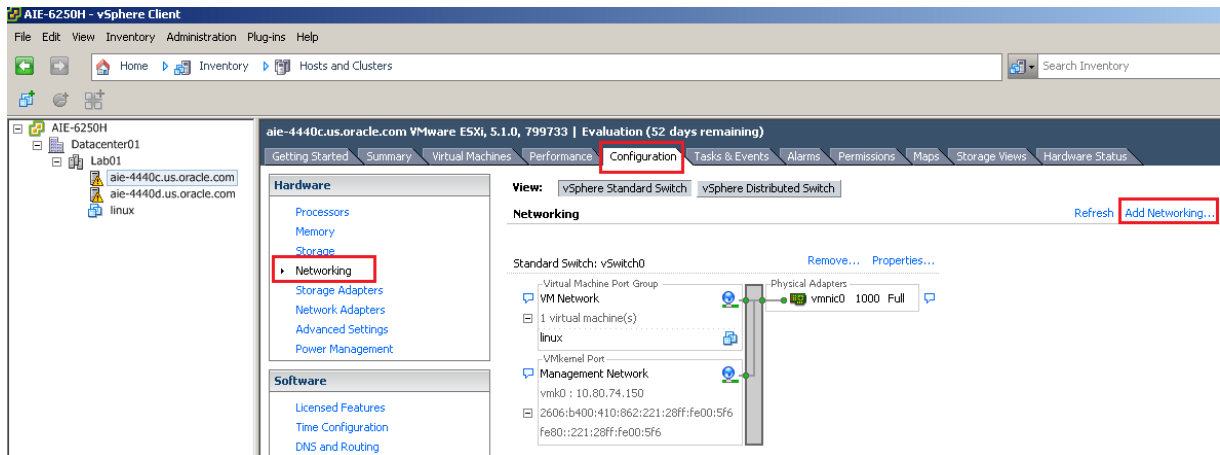


Figure 13. Configuring a VMware virtual switch shown on the VMware vSphere client

The windows for the remaining configuration steps are all displayed in the following figure.

2. In the next screen, in the connection type box, select **VMkernel** and click on Next.
3. In the VMkernel network access screen, select the network adapters that will be part of your new virtual switch. The example shows selection of the two InfiniBand network adapters, vmnic_ib0 and vmnic_ib1. Click on Next.
4. In the Port Group Properties section, enter a network label and select a VLAN ID that best matches your virtualized environment. Leave the 'Use the port group for vMotion,' 'Use the port group for fault tolerant logging' and 'Use the port group for management traffic' options unchecked. Click on Next.
5. In the next window, enter the IP address and subnet mask for your new network connection and click on Next.
6. Review your network configuration (this window is not seen in the example figure) and click on Finish to create your new VMware virtual switch.

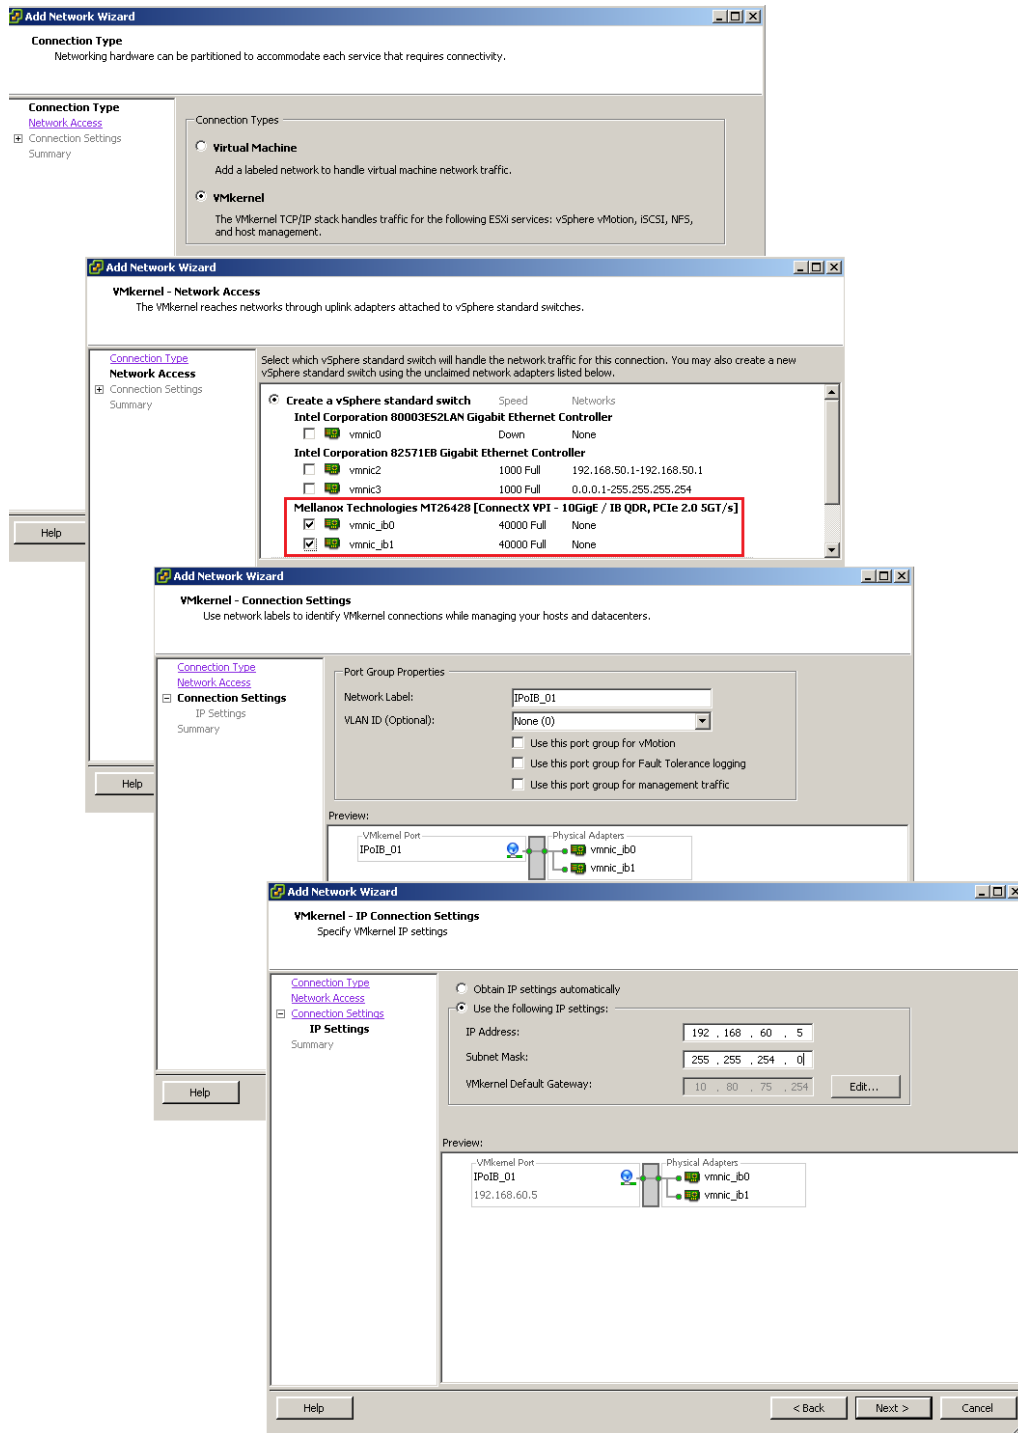


Figure 14. Step-by-step configuration for the VMware virtual switch in the VMware vSphere client

Once your virtual switch has been successfully configured, configure the network teaming, policy exception and network card failover order, as seen in the following figure and steps.

1. On the VMwre vSphere client, click on **Properties** of your new virtual switch, and then inside the Ports configuration, select **IPoIB** and click on Edit.
2. The IPoIB Properties screen will open. Select the **NIC Teaming** tab. Inside of Policy Exceptions, check the following:
 - Load balance: Route based on the originating virtual port ID
 - Network failover detection: Link status only
 - Notify switches: Yes
 - Failback: Yes
 - Failover Order: Select 'Override switch failover order.' Make both Infiniband over IP uplinks adapters active.

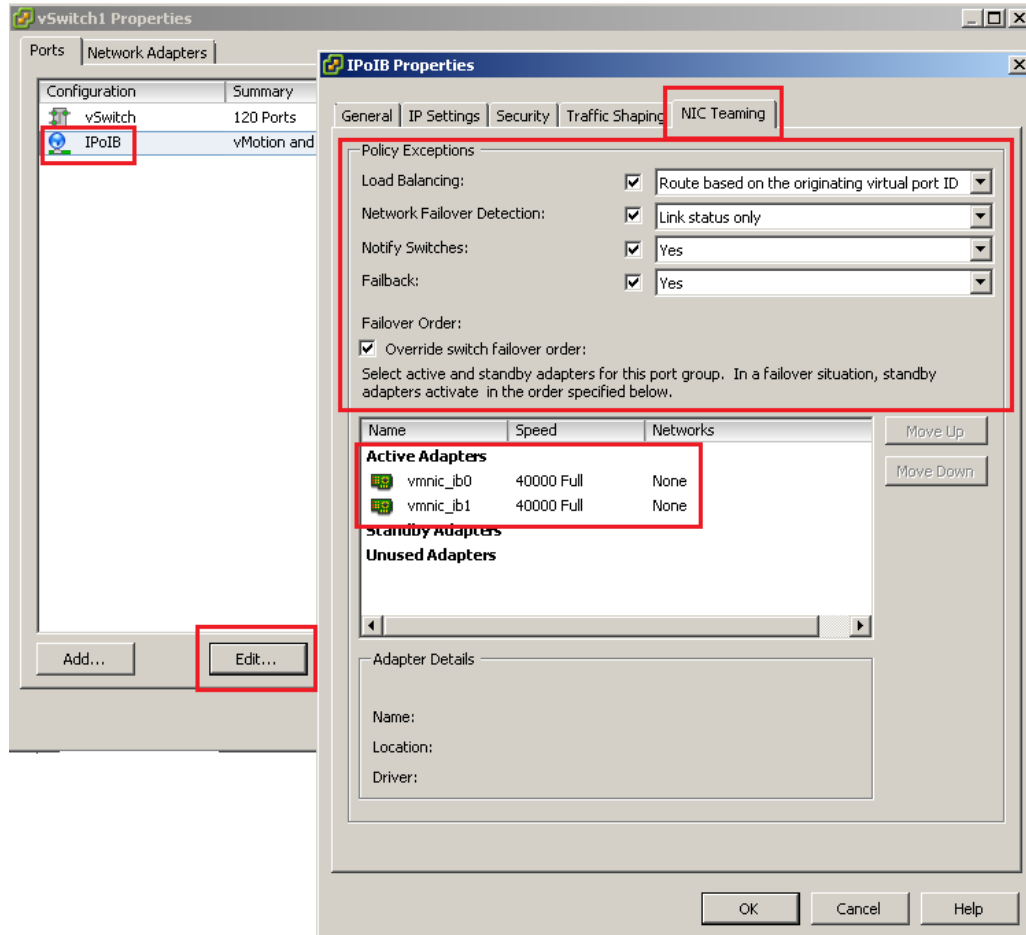


Figure 15. VMware virtual switch properties configuration

Recommendations for VMware ESXi5.1 Multipathing Settings for iSCSI

In this environment, the supported VMware Storage Array Type Plug-in (SATP) for Oracle ZFS Storage Appliance iSCSI LUNs is VMW_SATP_DEFAULT_AA and the supported VMware Path Selection Policy (PSP) is round-robin VMW_PSP_RR. It is important that you add the multipath policy before presenting or even scanning any new Oracle ZFS Storage Appliance LUNs.

The following VMware multipath policy needs to be included on your VMware ESXi5.1 hosts. To implement this, open an SSH session to your ESXi5.1 hosts and use the following instructions:

1. Identify if your ESXi5.1 hosts already have some policy for Oracle ZFS Storage Appliance using the following command. If the command does not return any value, there is no existing SATP storage policy for Oracle ZFS Storage Appliance and you can add the new rule.

```
esxcli storage nmp satp rule list | grep -i ZFS
```

2. Add the VMware multipath rule for Oracle ZFS LUNs.

```
esxcli storage nmp satp rule add --vendor="SUN" --model="^ZFS*" --satp=VMW_SATP_DEFAULT_AA --psp=VMW_PSP_RR
```

3. Ensure that the new rule has been successfully added so that you can present the Oracle ZFS Storage Appliance iSCSI LUNs.

```
esxcli storage nmp satp rule list | grep -i ZFS
```

```
VMW_SATP_DEFAULT_AA          SUN      ^ZFS*          user
VMW_PSP_RR
```

After presenting the iSCSI LUNs, ensure that the new SATP rule has claimed all new iSCSI LUNs to work with the right SATP and PSP multipath policies combination.

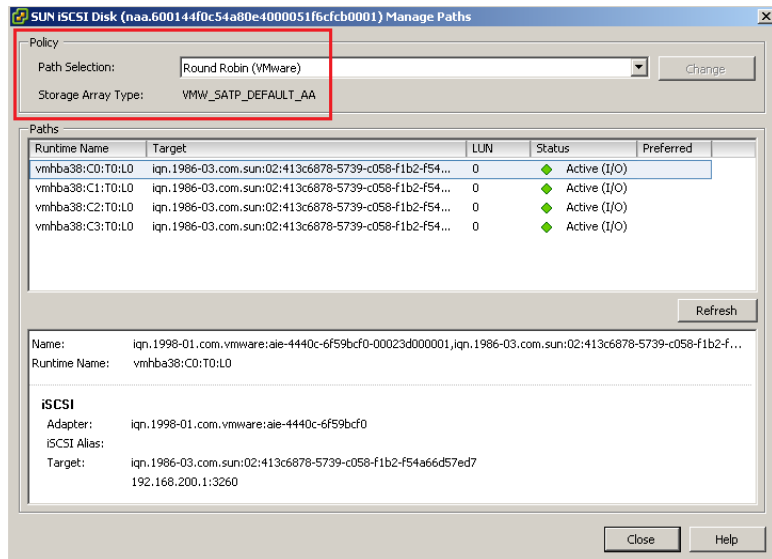


Figure 16. VMware PSP and SATP multipathing policy for iSCSI LUNs shown on VMware vSphere client

IMPORTANT: The configuration shown in the previous figure example is recommended for Oracle ZFS Storage Appliance iSCSI LUNs only and will not work for Fibre Channel LUNs. Once you add the rule shown, all Oracle ZFS Storage Appliance LUNs will be configured to work with the VMW_SATP_DEFAULT_AA policy and VMW_PSP_RR (round-robin) multipath policy combination. However, Oracle Fabric Interconnect provides different options for provisioning virtual host bus adapters (vHBAs) for your VMware environment, so you could potentially be working with both protocols (FC and iSCSI) at the same time and also in the same environment. To fix that, you will need to change the SATP policy for your Oracle ZFS Storage Appliance Fibre Channel LUNs, on which the VMW_SATP_ALUA policy will be applied.

In this environment there is no additional configuration for NFS protocol. You can use the same 40Gbps IPoIB for configuring VMware datastores based on NFS. However, ensure you have applied the NFS best practices, recommendations and tuning options as detailed in the following documentation, "How to Get the Best Performance from the Oracle ZFS Storage Appliance: Best Practices for Oracle ZFS Storage Appliance and VMware vSphere5.x."

<http://www.oracle.com/technetwork/articles/servers-storage-admin/sto-recommended-zfs-settings-1951715.html>

Apply the IOPs tuning options for all iSCSI and Fibre Channel LUNs. These options are also listed in the downloadable white paper version "Best Practices for Oracle ZFS Storage Appliance and VMware vSphere5.x":

<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/bestprac-zfssa-vsphere5-1940129.pdf>

Use the following steps to change the SATP multipath policy for your Fibre Channel LUNs only.

1. Identify all of the Fibre Channel LUNs attached to your ESXi5.1 host and for which the SATP multipath policy will change. The following example shows the Fibre Channel LUN naa.600144f0d230d6c5000051ffde250001.

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

```
naa.600144f0d230d6c5000051ffde250001
```

2. Alter the SATP multipath policy for this LUN and also enable the policy.

```
esxcli storage nmp satp rule add --device=
naa.600144f0d230d6c5000051ffde250001--satp=VMW_SATP_ALUA --
psp=VMW_PSP_RR
```

3. Reclaim the new configuration.

```
esxcli storage core claiming reclaim --
device=naa.600144f0d230d6c5000051ffde250001
```




Figure 17. VMware PSP and SATP multipathing policy for Fibre Channel LUNs shown on VMware vSphere Client

Conclusion

Oracle Fabric Interconnect's architecture and features for network virtualization were developed to provide fast throughput, low latency and an intelligent management, providing outstanding performance for virtualized environments. Oracle ZFS Storage Appliance offers an intelligent caching technology designed to deliver thousands of IOPS for your virtualized environment as well as the best throughput and response time for your virtualized applications and databases. In combination with VMware, Oracle ZFS Storage Appliance and Oracle Fabric Interconnect will deliver a high performance and low latency virtualized environment to address the demands of your critical virtualized application.

Appendix A: Updating Mellanox HCA Firmware

The following section details how to update a Mellanox ConnectX/ConnectX-2 HCA card in a VMware environment.

IMPORTANT: Before updating the firmware of your InfiniBand HCA, ensure that you have the correct firmware for your HCA model. Installation of wrong firmware can damage the HCA and make it unusable. Contact your HCA vendor to obtain the correct firmware for your card. Also, if you are working with an Oracle Infiniband HCA, refer to support.oracle.com for the latest HCA firmware release.

Ensure that your Mellanox ConnectX® or ConnectX®-2 are running with firmware version 2.9.1000 or higher, and the ConnectX®-3 version 2.10.0000 or higher.

The listed firmware releases are requirements to have ConnectX cards working with multifunction mode (10GbE and IB). For further information, please refer to the Mellanox web site.

Ensure that the HCA card is supported by your hardware (servers/hosts). Contact oracle.support.com for more information.

There are two ways to update the firmware of your InfiniBand HCA: the “push” option or a LINUX open source live CD. With the push option, you can use your Oracle Fabric Interconnect director to push the HCA firmware to your ESXi host, and then have your HCA firmware updated. With the live CD option you can work with the open source LINUX live CD distribution from which to boot your ESXi hosts, and then install the Oracle Virtual Networking utility for updating your HCA firmware.

To perform a “push” update, follow these steps:

1. Ensure that your ESXi5.1 hosts already have the Oracle Virtual Networking driver installed.
2. Identify the firmware release your HCA cards currently have. In the Oracle Fabric Interconnect CLI, run the following command:

```
show physical-server your_server_name hca
```

You will have an output similar to the following. The version row will tell you which firmware release your HCA cards are currently working with.

```
admin@qdr1[xsigo] show physical-server aie-6275m2a0
name      guid          descr      port      os          version      server-profile
-----
aie-6275m2a0  21280001efe456  qdr1:ServerPort12  VMware/ESXi-5.3.1.ESX.1/x86_64  2.11.2010/3.0.0
aie-6275m2a0  21280001efe455  qdr1:ServerPort13  VMware/ESXi-5.3.1.ESX.1/x86_64  2.11.2010/3.0.0
2 records displayed
admin@qdr1[xsigo]
```

Figure 18. Identifying the HCA cards' firmware release

3. Using an SSH client, copy the right firmware for your HCA cards to the /tmp of your Oracle Fabric Interconnect switch.
4. On the Oracle Fabric Interconnect CLI, run the following command line to push the new firmware to your ESXi 5.1 hosts.

```
set physical-server your_server_name upgrade-hca your_new_firmware_name
```

To perform the 'Live CD' update, follow these steps:

1. Boot your ESXi5.1 hosts with the Linux open source live CD.
2. Configure the network of your new Linux live environment.
3. Download the [Oracle Virtual Networking Tool to burn option ROM 4.0.3 for Linux](#) and the correct firmware to your HCA cards. Using an SSH client, copy both to your Linux live environment.
4. Unzip and install the .RPM of Oracle Virtual Networking Tool to burn option ROM. The CLI session should look similar to the following:

```
[root@localhost ~]# cd /tmp/
[root@localhost tmp]# ls -ltr
total 21636
-rw-r--r--. 1 root root 21488017 Jul  9 00:24 xsigo-hca-firmware-
4.0.3.r177.XGBOOT-1.i386.zip
-rw-r--r--. 1 root root 661228 Jul  9 12:24 SUN0170000009_fw.img
-rw-----. 1 root root 0 Jul 25 06:27 yum.log

[root@localhost tmp]# unzip xsigo-hca-firmware-4.0.3.r177.XGBOOT-1.i386.zip
Archive:  xsigo-hca-firmware-4.0.3.r177.XGBOOT-1.i386.zip
  inflating: xsigo-hca-firmware-4.0.3.r177.XGBOOT-1.i386.rpm

[root@localhost tmp]# rpm -ivh xsigo-hca-firmware-4.0.3.r177.XGBOOT-
1.i386.rpm
Preparing...          #####
[100%]
   1:xsigo-hca-firmware  #####
[100%]

[root@localhost tmp]# /opt/xsigo/bin/xg_config

#####
###
# (C) 2007,2008 XSIGO SYSTEMS Inc. All rights reserved. This material may not
be
# reproduced, displayed, modified or distributed without the express prior
# written permission of the copyright holder.
#####
###

#####
###
# Main menu
#####
###

Selected card:
Node GUID      : '0021280001efe454'
Board ID      : 'SUN0170000009'
```

```

CA type           : 'InfiniBand: Mellanox Technologies MT26428 [ConnectX
VPI PCIe 2.0 5GT/s - IB QDR / 10GigE] (rev b0) '
Firmware version  : '2.7.8130'
Hardware version  : 'b0'
Option ROM version : 'unknown'

```

```

1) Flash HCA Firmware
2) Burn Option ROM
3) Erase Option ROM
4) Change selected card
0) Quit
Select option> 1

```

```

#####
###
# Flash HCA Firmware Menu
#####
###

```

```

Selected card:
Node GUID       : '0021280001efe454'
Board ID        : 'SUN0170000009'
CA type         : 'InfiniBand: Mellanox Technologies MT26428 [ConnectX
VPI PCIe 2.0 5GT/s - IB QDR / 10GigE] (rev b0) '
Firmware version : '2.7.8130'
Hardware version  : 'b0'
Option ROM version : 'unknown'

```

Press any key to continue

Note: If you get the error message "Found no matching firmware files to use," use the `mstflint` utility to manually update your HCA firmware with the following steps:

1. On the same Linux live environment, identify the Mellanox card using the following command (determines the PCI bus ID):

```
lspci | grep -i "mellanox"
```

```

00:03:00.0 Serial bus controller: Mellanox Technologies MT26428
[ConnectX VPI PCIe 2.0 5GT/s - IB QDR / 10GigE]

```

2. Using an SSH client, copy the correct HCA firmware to your Linux live environment, and unzip the firmware image by running `unzip firmware file name.zip`. Use the `mstflint` utility for updating the firmware of your HCA card.
3. In the same directory as the firmware file, run:

```
/opt/xsigo/bin/mstflint -d <PCI bus ID> -i <.bin file> b".
```

```

[root@localhost tmp]# /opt/xsigo/bin/mstflint -d 03:00.0 -i fw-
ConnectX2-rel-2_11_2010-be-375-3696-01.bin b

```

```
Current FW version on flash: 2.7.8130
```

New FW version: 2.11.2010

Burning FW image without signatures - OK

Restoring signature - OK

For more information, refer to Mellanox firmware update documentation and also My Oracle Support at support.oracle.com.

http://www.mellanox.com/page/firmware_HCA_FW_update

For details on `mstflint` utility usage, see [mstflint FW Burning Tool README](#).

4. Run `ibstat` to verify the new firmware version.
5. Reboot the system without the Linux live CD. Once you have your ESXi5.1 host up and running again, verify the new firmware version using the command line `show physical-server your_server_name hca` on the Oracle Fabric Interconnect CLI.

Appendix B: References

Oracle ZFS Storage Appliance Documentation

RESOURCES	LOCATION
<i>Sun ZFS Storage Appliance Administration Guide</i>	Oracle ZFS Storage Appliance documentation library at
<i>Sun ZFS Storage 7000 Analytics Guide</i>	http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html
<i>Sun ZFS Storage 7x20 Appliance Installation Guide</i>	
<i>Sun ZFS Storage 7x20 Appliance Customer Service Manual</i>	
VMware	http://www.vmware.com
Oracle Fabric Interconnect Documentation	http://docs.oracle.com/cd/E38500_01/
"Configuring Xsigo Systems I/O Directors on ESX 4.x, ESXi 4.x and ESXi 5.x (Partner Verified and Support) (2005239)"	http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2005239
VMware vSphere 5.1 Documentation	http://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-pubs.html
Mellanox	http://www.mellanox.com



VMware vSphere 5.x with Oracle® ZFS Storage
Appliance and Oracle Virtual Networking
October 2013, Version 1.0
Author: Anderson Souza

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



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

Copyright © 2013, 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 licensed through X/Open Company, Ltd. 0611

Hardware and Software, Engineered to Work Together