

Implementing Fibre Channel SAN Boot with the Oracle ZFS Storage Appliance

January 2014

By Tom Hanvey; update by Peter Brouwer

Version: 2.0

This paper describes how to implement a Fibre Channel (FC) SAN boot solution using Oracle ZFS Storage Appliance on a high-availability SAN. This solution has been validated with a variety of servers, operating systems, and hardware configurations.

[Introduction](#)

[Overview](#)

[Benefits of Booting from a SAN](#)

[FC SAN Boot Configuration Requirements](#)

[Configuring the Client Server for Fibre Channel SAN Boot](#)

[Installing the Operating System on the Server](#)

[Best Practices](#)

[Conclusion](#)

[Appendix: References](#)

Introduction

Organizations are continually looking for ways to simplify their management infrastructure, enhance scalability, and reduce costs while increasing reliability and performance in their data centers. Booting from a storage area network (SAN) offers many benefits, leading to cost savings as well as higher levels of protection, ease of management, increased flexibility, and reduced down time.

Traditionally, operating systems have been installed on internal disks on individual servers or on direct attached storage (DAS). This approach presents a number of challenges to an IT organization. Dedicated internal boot devices cannot be shared with other servers, so are often underutilized. IT staff must perform management tasks on these systems locally rather than from a central management system, leading to increased administrative costs. For optimal redundancy and performance, additional RAID software or host bus adapters (HBAs) are required to manage these storage devices.

Local disks on individual servers present particular challenges for multisite administration and disaster recovery site maintenance. Creating clones of disk content on off-site hosts or replicating server operating systems to a disaster recovery backup site can be complex operations requiring specialized software.

The complex task of managing the servers of an entire enterprise can be simplified by enabling data center administrators to centrally manage all storage-related tasks, such as operating system maintenance, at the array level rather than at the individual server level. Locating server boot devices on an Oracle ZFS Storage Appliance accessed by servers across a high-availability Fibre Channel (FC) SAN enables increased efficiency, and even automation, of many administrative tasks, significantly reducing operating expenses.

If a server goes down, a system administrator can boot up a standby server in a matter of minutes to resume business. Snapshots and clones of operating system images stored on the Oracle ZFS Storage Appliance can be simultaneously deployed to servers in development and test environments or to secondary disaster recovery sites.

Booting from the SAN reduces the amount of time required for server upgrades. With minimal reconfiguration, you can replace an outdated or underpowered server with a new server, which you can then point to the original FC SAN boot device.

By locating server boot devices on a RAID-protected shared storage device like the Oracle ZFS Storage Appliance, you can eliminate the need for hardware or software RAID devices in each server, which helps reduce hardware costs.

Overview

A boot-from-SAN solution implemented with an Oracle ZFS Storage Appliance located on a high-availability FC SAN is shown in Figure 1. In this solution, servers are booted from a pool of centrally managed storage volumes located in the Oracle ZFS Storage Appliance. Each storage volume in the pool serves as a boot LUN for a specific server. The figure shows two types of servers used in the validation testing (x86, and Solaris SPARC) and the operating systems validated on each server type. Note that each server platform is also available in a modular hardware blade architecture.

When the Oracle ZFS Storage Appliance is also used for data storage, best practices dictate that a dedicated pool and separate data paths be used for boot devices. See the following section "Best Practices" for more details.

You can use any server and host bus adapter (HBA) combination that supports SAN boot to implement an FC boot solution using an Oracle ZFS Storage Appliance.




Figure 1. Fibre Channel boot solution using an Oracle ZFS Storage Appliance

Benefits of Booting from a SAN

A boot from FC SAN solution provides significant benefits:

- **Reduces data center footprint and facility costs** - Booting from FC SAN enables you to use diskless servers and blade servers, which take up less space, consume less power, and require less cooling.
- **Lowers administrative overhead** - All operating system storage is provisioned and managed from the Oracle ZFS Storage Appliance. A server can be easily replaced by re-mapping its corresponding boot LUN to a new server. If the new server has the same profile as the server being replaced, it will boot the operating system from the SAN without requiring reconfiguration. Snapshots and clones of operating system images can be created and mapped to new servers on the SAN with just a few clicks of a mouse, simplifying migration and scalability tasks.
- **Facilitates disaster and server failure recovery** - By installing operating systems on the Oracle ZFS Storage Appliance rather than individual servers, you can take advantage of the data protection and redundancy features of the Oracle ZFS Storage Appliance to help reduce downtime during maintenance and fault outages. Operating system images can be protected using snapshots and clones or backed up using Network Data Management Protocol (NDMP).

FC SAN Boot Configuration Requirements

Configuring a Fibre Channel boot solution using an Oracle ZFS Storage Appliance requires the following:

- Zoning must be configured in the local SAN such that the server FC ports can see the Oracle ZFS Storage FC target ports. For more details, refer to documentation in Appendix A: References.

- In the Oracle ZFS Storage Appliance, at least one FC PCIe card must be installed with one port enabled in target mode. For details, see the section "Configuring the Oracle ZFS Storage Appliance for Fibre Channel Boot" that follows.
- An HBA that supports SAN boots must be installed in each server to be provisioned from the SAN. The solution described in this paper was tested with the following FC HBA driver and firmware versions:
 - QLogic QLE2562 (Firmware Version 4.03.02, BIOS Revision 2.02)
 - Emulex LPe12002 (BIOS Version 2.11a2)
 - The FC HBA on each server must be configured as the primary boot device, a storage target LUN in the Oracle ZFS Storage Appliance must be provisioned with the appropriate operating system, and the LUN mapped to the server's initiator port. For details, see the section "Configuring the Client Server for Fibre Channel SAN Boot" that follows.

Configuring the Oracle ZFS Storage Appliance for Fibre Channel Boot

To configure the Oracle ZFS Storage Appliance for FC SAN boot, complete the following steps:

1. Check that at least one FC PCIe card is installed in the Oracle ZFS Storage Appliance.
2. By default, all the FC ports on the Oracle ZFS Storage Appliance are set to initiator mode. To enable a port in target mode, complete these steps:
 - a. Log in to the Oracle ZFS Storage Appliance and navigate to Configuration > SAN > Fibre Channel Ports.
 - b. Set the selected port to Target mode as shown in Figure 2.
 - c. Click the Apply button.

NOTE: Changing this setting requires a reboot.



Figure 2. Setting a PCIe port to target mode in the Oracle ZFS Storage Appliance

3. Provision each LUN that will serve as a server boot LUN with the appropriate initiator and target groups.

Configuring the Client Server for Fibre Channel SAN Boot

To configure each client server for an FC SAN boot, first confirm that a Fibre Channel HBA is installed on the client and that the HBA supports SAN boot. Then set the boot precedence in the system BIOS to make the FC HBA card the highest priority boot device and configure the HBA to boot from the LUN on which the operating system for that server has been installed in the Oracle ZFS Storage Appliance. These procedures are described in the following sections.

Setting Boot Precedence in the System BIOS

Set the boot precedence in the system BIOS so that the FC HBA card is the highest priority boot device by completing these steps:

1. Reboot the server. While the server is initializing, press F2 to display the system BIOS menu.
2. If the server has an LSI PCI card installed, disable the PCI slot in the system BIOS. In some servers, such as Oracle's Sun x86 servers, the LSI card takes higher boot precedence and will try to boot the local operating system. To prevent this from happening, complete the following steps:
 - a. Select Advanced to display the PCI Configuration screen.
 - b. Disable the PCI slot in which the LSI card is installed as shown in Figure 3.
 - c. Press F10 to save the setting and exit the screen and reboot the server.

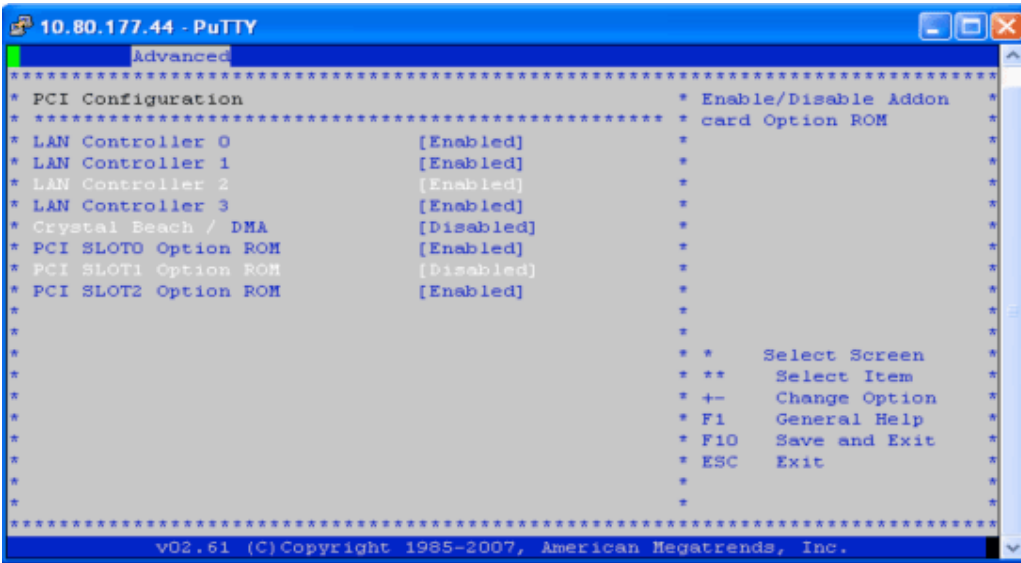


Figure 3. System BIOS PCI Configuration screen showing PCI Slot1 disabled

3. Set the FC HBA card to be the highest priority boot device.
 - a. From the BIOS menu, select Boot to display the Boot Device Priority screen.
 - b. Select the FC HBA as the 1st Boot Device as shown in Figure 4.
 - c. Press F10 to save settings, exit, and reboot the server.

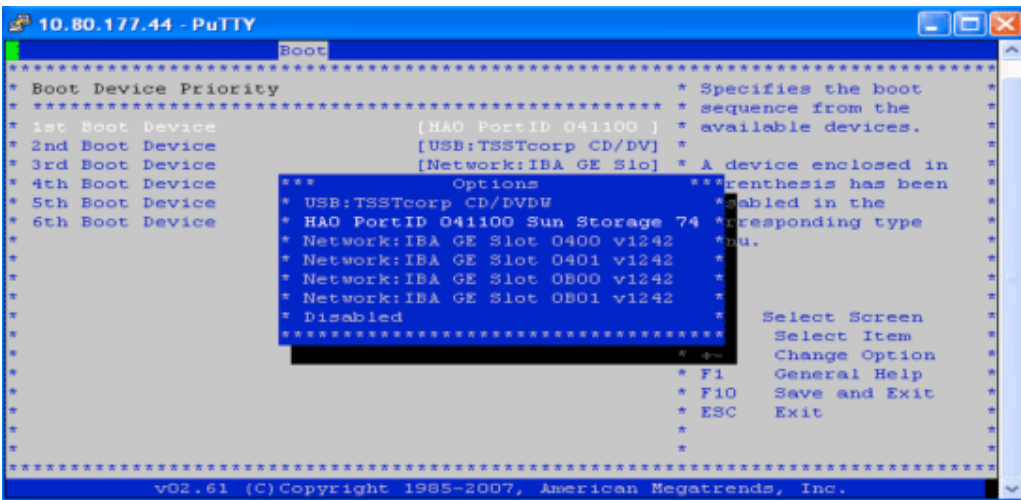


Figure 4. System BIOS PCI Configuration screen showing the FC HBA set as the primary boot device

Configuring the Host Bus Adapter for Fibre Channel Boot

One or more ports on the FC HBA on the server must be configured to boot from the LUN on which the operating system for that server has been installed in the Oracle ZFS Storage Appliance. The following procedure shows the steps for a QLogic FC HBA. The procedure is similar for an Emulex FC HBA.

1. Reboot the system. When the initialization screen shown in Figure 5 appears, log in to the HBA BIOS menu.


```

10.80.177.44 - PuTTY
*****Selected Adapter*****
*****
* Adapter Type Address Slot Bus Device Function *
*****
* QLE2562 7C00 28 03 00 0 *
*****

*****Configuration Settings*****
*****
* Adapter Settings *
* Selectable Boot Settings *
* Restore Default Settings *
* Raw Nvram Data *
* Advanced Adapter Settings *
*****

Use <Arrow keys> to move cursor, <Enter> to select option, <Esc> to backup

```

Figure 8. Accessing the Adapter Settings for the HBA

5. On the Adapter Settings screen, select Host Adapter BIOS and press ENTER to enable the host adapter BIOS as shown in Figure 9 (the host adapter BIOS is disabled by default).

```

10.80.177.44 - PuTTY
*****Selected Adapter*****
*****
* Adapter Type Address Slot Bus Device Function *
*****
* QLE2562 7C00 28 03 00 0 *
*****

*****Adapter Settings*****
*****
* BIOS Address: C6800 *
* BIOS Revision: 2.10 *
* Adapter Serial Number: L86194 *
* Interrupt Level: 10 *
* Adapter Port Name: 2100001B329219 *
* Host Adapter BIOS: Enabled *
* Frame Size: 2048 *
* Loop Reset Delay: 5 *
* Adapter Hard Loop ID: Disabled *
* Hard Loop ID: 0 *
* Spinup Delay: Disabled *
* Connection Options: 2 *
* Fibre Channel Tape Support:Enabled *
* Data Rate: 2 *
*****

Use <Arrow keys> and <Enter> to change settings, <Esc> to exit

```

Figure 9. Enabling the host adapter BIOS

6. To change the boot device priority level, press <Esc> to return to the Configuration Settings menu and select Selectable Boot Settings as shown in Figure 10. A list of available FC target ports displays, as shown in Figure 11.

```

10.80.177.44 - PuTTY
*****Selected Adapter*****
*****
* Adapter Type Address Slot Bus Device Function *
*****
* QLE2562 7C00 28 03 00 0 *
*****

*****Configuration Settings*****
*****
* Adapter Settings *
* Selectable Boot Settings *
* Restore Default Settings *
* Raw Nvram Data *
* Advanced Adapter Settings *
*****

Use <Arrow keys> to move cursor, <Enter> to select option, <Esc> to backup

```

Figure 10. Accessing the HBA boot settings

7. Select the FC target port the HBA will use on the Oracle ZFS Storage Appliance as shown in Figure 11 and press <Enter>. The Selectable Boot Setting screen for the HBA port displays as shown in Figure 12.




Figure 11. Selecting the FC target port the HBA will use on the Oracle ZFS Storage Appliance

8. Select (Primary) Boot Port Name as shown in Figure 12 and press <Enter> to display a list of all the available LUNs as shown in Figure 13.

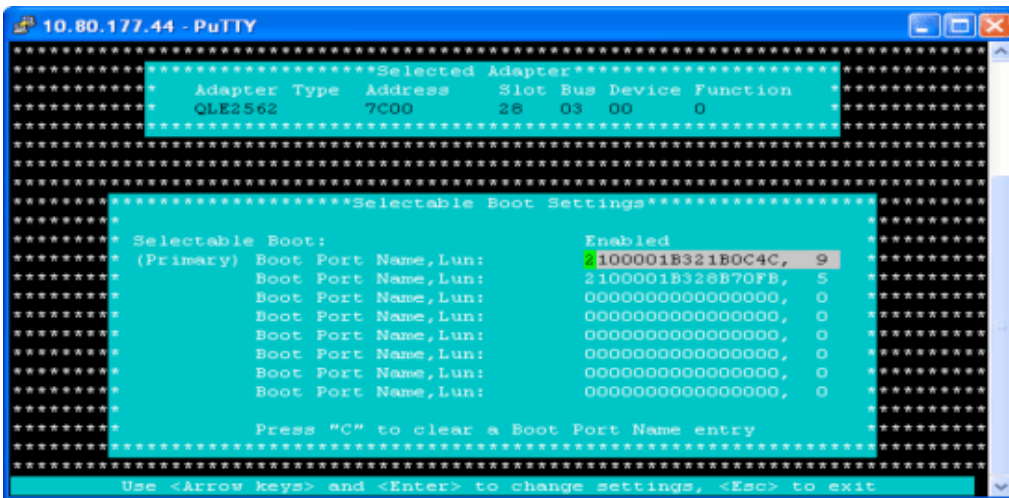


Figure 12. Selecting the primary boot port

9. Select the number of the LUN from which the server operating system should boot as shown in Figure 13.

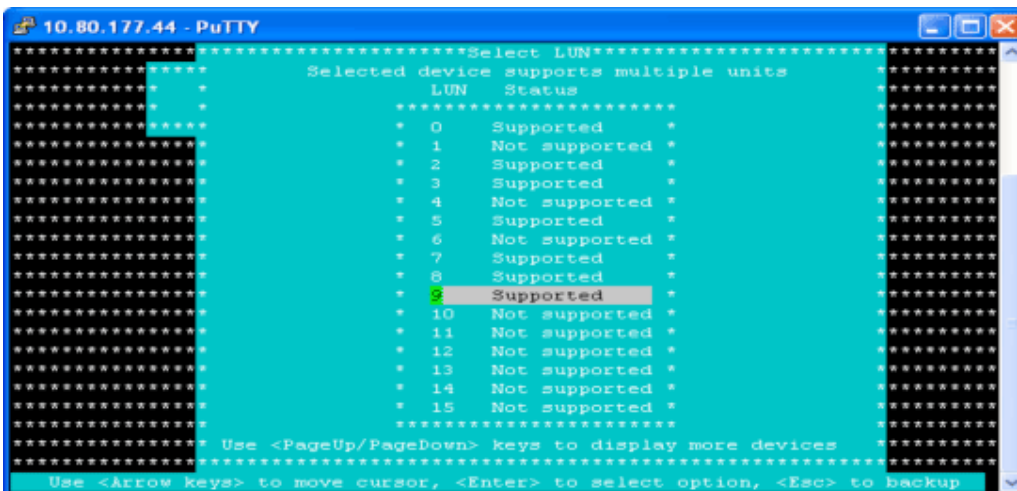


Figure 13. Selecting the boot LUN for the server

When configuring the Emulex BIOS, you have the option to boot the server using the World Wide Port Name (WWPN) or Device ID as shown in Figure 14.

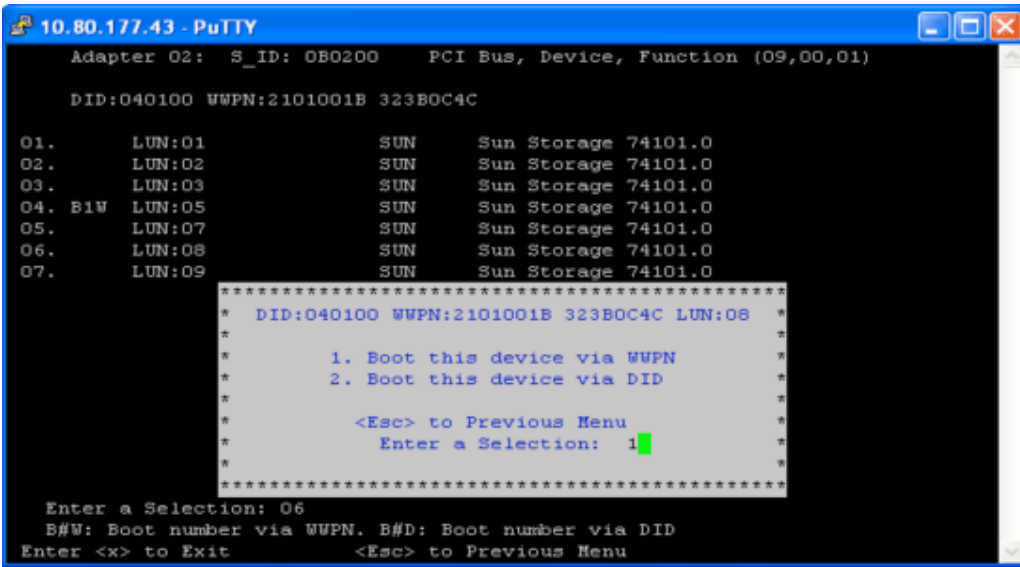


Figure 14. Selecting the boot device identification method when configuring an Emulex HBA

10. For a second HBA port, repeat Steps 2 through 10. Use the same settings as for the first HBA port.
11. Press <Esc> and save the configuration settings as shown in Figure 15.

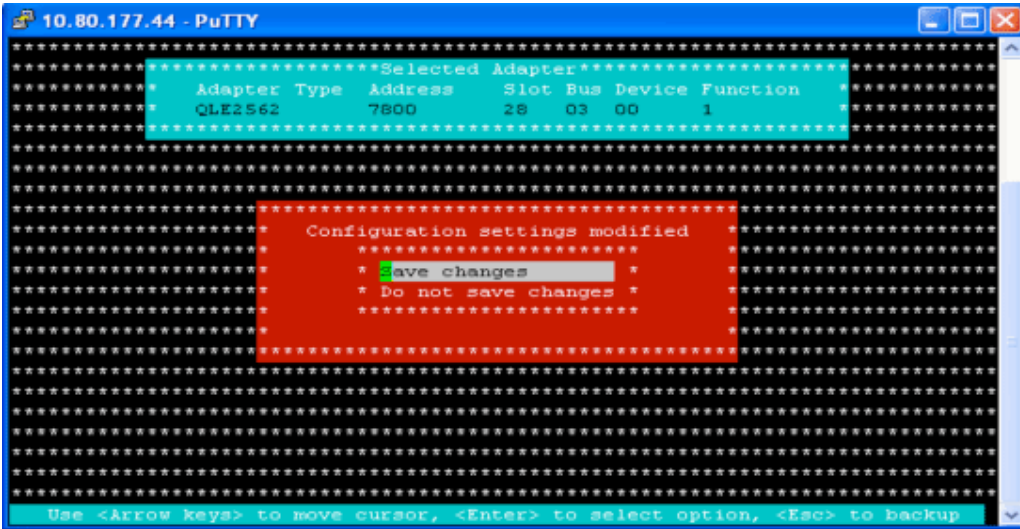


Figure 15. Saving the HBA configuration settings

12. Reboot the server.

When the server boots, it will now choose the FC HBA as the primary boot device. It will use the primary boot setting in the HBA BIOS to select the appropriate LUN in the Oracle ZFS Storage Appliance from which to boot the operating system.

Installing the Operating System on the Server

To install the operating system on the server, follow the instructions for the specific operating system as shown.

Installing Oracle Solaris

To install Oracle Solaris on a server, during the installation process, select the appropriate FC LUN from which to install the operating system.

Installing Oracle Linux

To install Oracle Linux (5 u4) on a server, during the installation process, select Advanced Configuration to install the GRUB boot loader and the operating system on the same FC LUN device. Otherwise, the GRUB master boot record (MBR) will be installed on the local disk and the operating system will not boot from the primary boot FC LUN.

Installing SUSE Linux Enterprise 11

To install SUSE Linux Enterprise 11 on a server, during the installation process, select Advanced Configuration to install the GRUB boot loader and the operating system on the same FC LUN device. Otherwise, the GRUB master boot record (MBR) will be installed on the local disk and the operating system will not boot from the primary boot FC LUN.

Installing Microsoft Windows 2003

To install Microsoft Windows 2003 on a server, complete the following steps:

1. Create an image on a floppy disk of the driver for the QLogic or Emulex FC HBA installed on the server.
2. Reboot the server to start the installation process. During installation initialization, press F6 to provide the path to the HBA driver image on the floppy disk.
3. Continue the installation using the HBA driver image. The driver enables the HBA to see the FC LUN that is serving as the primary boot LUN for the operating system installation.
4. Install the operating system on the new FC bootable LUN.

Installing Microsoft Windows 2008

To install Microsoft Windows 2008 on a server, complete the following steps:

1. Set the boot LUN as the primary boot device on the FC HBA on the server.
2. Proceed with the installation.

Best Practices

Consider these practices for optimal configuration of the Oracle ZFS Storage Appliance:

Configure storage on the Oracle ZFS Storage Appliance for the highest level of fault tolerance for FC SAN boot solutions. Storage pool(s) should be configured for optimum fault tolerance and performance by using mirroring.

Storage LUNs should be mapped to multiple HBA ports accessing the SAN. Identify multiple ports on the client server in the systems BIOS. The BIOS will go through the list of targets upon boot until it can find the active path to the boot device. NOTE: Asynchronous Logical Unit Access (ALUA) is not supported for boot solutions because most HBAs do not support ALUA in the HBA firmware.

When configuring an FC boot device in an Oracle ZFS Storage Appliance that is also hosting LUNS for other applications, be sure to separate the boot paths from the application data paths. Sharing FC target ports and storage pools designated for booting servers with FC target ports and storage pools servicing applications can adversely affect performance and is not recommended.

Configure dedicated boot storage pools (mirrored) and separate application pools (variable).

Map targets across alternate target ports on the Oracle ZFS Storage Appliance so they are not shared with boot ports.

Conclusion

A Fibre Channel SAN boot solution using the Oracle ZFS Storage Appliances on a high availability SAN provides a high level of protection while lowering administrative overhead. Your operating system assets can be protected using a variety of methods

including snapshots, clones, replication, or Network Data Management Protocol (NDMP) backup to tape. The ability to manage and maintain operating systems on servers throughout the organization from a central appliance lowers administrative costs.

Appendix: References

For troubleshooting information related to setting up the FC driver and ALUA, see: SAN Configuration and Multipathing topic, Oracle Solaris documentation

<http://docs.oracle.com/cd/E19082-01/820-3070/index.html>

Note: The ALUA path failover feature used in the Oracle ZFS Storage Appliance cluster configurations is not supported for boot devices.

Other useful links:

Oracle ZFS Storage Appliance documentation; see Chapter 3, Configuration, in the Administration Guide
<http://www.oracle.com/technetwork/documentation/oracle-unified-ss-193371.html>

SUSE Linux Enterprise Server 11 documentation
http://www.suse.com/documentation/sles11/stor_admin/?page=/documentation/sles11/stor_admin/da/ta/be5rvii.html

"Multipathing Support in Windows Server 2008," blog
<http://blogs.msdn.com/b/san/archive/2008/07/27/multipathing-support-in-windows-server-2008.aspx>

Enterprise Linux Multipathing
https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html-single/DM_Multipath/



Implementing Fibre Channel SAN Boot with the
Oracle ZFS Storage Appliance
January 2014, Version 2.0
Authors: Application Integration Engineering,
Tom Hanvey; update by Peter Brouwer

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 © 2011, 2014 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