ORACLE

ZFS STORAGE APPLIANCE

An Oracle Technical White Paper
January 2014

Implementing iSCSI Boot with Oracle Solaris on the Oracle ZFS Storage Appliance

ORACLE

Table of Contents

Table of Figures

Introduction

The iSCSI boot feature of Oracle Solaris enables you to load and start Oracle Solaris over the network from a remote location. The iSCSI boot feature supports booting from both SPARC-based and x86-based hardware systems. iSCSI (Internet Small Computer System Interface) boot is typically initiated by the server while the boot image resides on an iSCSI target that is attached to the network. Because the feature uses a standard Ethernet-based infrastructure, only standard network infrastructure components are required.

This paper describes how to set up the Oracle ZFS Storage Appliance to use iSCSI LUNs from which to boot Oracle Solaris, how to install Oracle Solaris on the configured iSCSI LUN, and how to configure the server boot environment to boot from the installed Oracle Solaris image.

This paper also details the use of Dynamic Host Configuration Protocol (DHCP) services. Using DHCP simplifies the Oracle Solaris installation process and it can be configured to supply the iSCSI boot information needed by the server's hardware boot process.

Note: In this paper a Network Interface Adapter (NIC) of the Intel 1000/Pro server series is used for the x86-hardware-based examples.

Benefits of Booting Using iSCSI

An iSCSI boot solution provides significant benefits:

- Reduces data center footprint and facility costs – Booting using iSCSI 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 centrally provisioned and managed from the Oracle ZFS Storage Appliance. A server can be easily replaced by remapping its corresponding boot LUN to a new server. If the new server has the same hardware profile as the server being replaced, the server will boot the operating system from the iSCSI LUN 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 normally caused by maintenance and fault outages. Operating system images can be protected using snapshots and clones, backed up using Network Data Management Protocol (NDMP) or images replicated to a second ZFS Storage Appliance on an off-site location.

Differences from a Network PXE Boot

Using iSCSI boot on a SPARC-based system to boot over the network differs from a typical SPARC network PXE (Pre-boot eXecution Environment) boot by its streamlined process. The network boot process is a multistep process; it requires loading a boot environment from a server using a TFTP (Trivial File Transfer Protocol) service, followed

by loading the initial Oracle Solaris boot loader. In comparison, the iSCSI boot environment loads an Oracle Solaris boot image directly from the iSCSI LUN.

Using iSCSI boot on an x86-based system to boot over the network differs from a typical x86 network boot in the following ways:

- A GRUB (Grand Unified Bootloader)-based network boot requires a DHCP server that is configured for PXE clients, whereas iSCSI boot does not.
- A PXE boot requires a boot server to provide the ramdisk image, whereas iSCSI boot does not.

For more information about booting a system from the network, see "x86: How to Perform a GRUB Based Boot From the Network" in the Oracle Solaris System Administration Guide at: http://docs.oracle.com/cd/E19082-01/819-2379/fvzpk/index.html.

iSCSI Boot Configuration Requirements

iSCSI boot has the following configuration requirements. An easily printed flowchart showing the overall installation process, including how and when these configuration requirements are implemented in that process, is in Appendix C: Configuration Information Handling near the end of this paper.

Oracle Solaris SPARC Platforms

iSCSI boot on SPARC platforms is supported with OpenBoot level 4.31 or later, and does not require a specific Network Interface Card (NIC). The boot command in OpenBoot takes a series of keywords to identify the destination iSCSI target or uses the parameters stored in the network-boot-parameters NVRAM variable. The command uses the format `boot net:`*`keyword=value`*.

x86 Platforms

On x86 platforms, the host that is being booted must use NICs that are iSCSI Boot Firmware Table (iBFT) capable or have a main board BIOS (Basic Input-Output System) that is iBFT capable. To configure iSCSI boot properly, refer to the documentation for your specific NIC hardware.

Oracle Solaris Operating Systems

iSCSI boot is supported as of Oracle Solaris release 11.1. During the installation process the iSCSI option leads to a screen that collects all information to install and set up Oracle Solaris for an iSCSI boot environment. This installation step requires a working IP connection at this point. The Oracle Solaris boot image requires access to a DHCP service for it to provide the necessary information to configure a network connection. When a DHCP service is not available, a network connection can be set up using the shell in the installation process. See Appendix B: DHCP Setup Tips.

Network Configuration and iSCSI Protocol Parameters

To successfully set up a connection at boot time between a server and an iSCSI LUN over a network, the boot program needs some network configuration information and some iSCSI protocol parameters.

Network information consists of server and target subsystem IP addresses, netmask and gateway information.

To set up an iSCSI connection between the server and the LUN subsystem requires the following information: an iSCSI initiator IQN (iSCSI Qualified Name) name, an iSCSI target IQN name, a LUN number, and the port number used by the LUN subsystem to listen for incoming iSCSI connection requests.

Additional information is needed when using authentication to gain access to the target LUN: either CHAP (Challenge Handshake Authentication Protocol) name and password login credentials or RADIUS (Remote Authentication Dial-In User Service) Server name and RADIUS login password.

The boot program must be able to gain access to this information. There are two options: the program either uses a DHCP server to obtain the information or it retrieves it from the local hardware NVRAM (Non-Volatile RAM).

The decision to use either NVRAM information or DHCP depends on user security requirements and boot hardware capabilities. A DHCP service cannot supply CHAP nor RADIUS authentication information. So the DHCP option cannot be used when login authentication is required.

When using a DHCP service, you can configure it to supply all network information, the target subsystem's IP address, IQN target name, LUN number and iSCSI listen port. Some DHCP implementations can support the handling of an iSCSI initiator IQN name. However, at the moment of this paper's writing, no boot hardware supports the retrieval of this information from a DHCP service.

Controlling the Boot Process on a SPARC Server

On a SPARC server, the boot process is handled by OpenBoot PROM (OBP). The boot program can either retrieve the configuration information from its NVRAM or use DHCP.

Using the command `boot:net[,`*`keyname1=value,keyname2=value,…`*`]` allows the boot program to use the information stored in the NVRAM network-boot-arguments variable. This variable also supports values for CHAP name and password.

Using the command `boot:dhcp` will let the boot program retrieve boot information from a DHCP service. It is currently not able to retrieve the iSCSI initiator IQN name.

When the boot program cannot retrieve the iSCSI initiator IQN name, it will use the following format: `iqn.1986-03.com.sun:boot.[NIC MAC in HEX].`

Controlling the Boot Process on an x86-based Server

The boot process on an x86 server is initially under control of the BIOS of the hardware. During boot, BIOS checks for whether Host Bus Adapters (HBAs) have any BIOS extension available. For support of iSCSI boot, NIC vendors supply a boot and configuration program that is stored on Flash memory of the NIC. During system boot the user can start the configuration program. This program lets the user enter the iSCSI

configuration information and stores it in NVRAM on the NIC. Depending on the NIC vendor specifications, a combination of NVRAM and DHCP supplied information can be used by the NIC iSCSI boot program. Check your vendor NIC documentation for further details.

DHCP Configuration Considerations

As mentioned earlier, DHCP can be used to retrieve the network information and some basic iSCSI setup information.

The Oracle Solaris installation procedure assumes availability of a DHCP service to set up a working network connection that will be used to discover the iSCSI LUN(s) upon which to install Oracle Solaris. If a DHCP service is not available, the network must be configured manually.

The question becomes whether the DHCP service should return a dynamic IP address or a fixed allocated address for this server. The answer depends upon whether the DHCP server will be used during the boot phase of the server too.

If no DHCP service will be used in the boot phase, then the server's IP address must be stored in NVRAM and will be used each time the server boots. In a SPARC environment, the Oracle Solaris installation process stores all required information in NVRAM. On an x86 platform, you must manually supply this information, so the DHCP service should return a fixed IP address that will be used permanently for this server.

When the DHCP service will be used during boot phase and the boot program supports the use of DHCP to retrieve the server's IP address, a dynamic IP address can be used.

As previously mentioned, there is currently no boot environment that supports the retrieval of the iSCSI initiator IQN value. For SPARC platforms you can make use of the fact that OBP uses a fixed format; use this format to configure the Oracle ZFS Storage Appliance iSCSI initiator IQN declaration as described in the next section. For x86 platforms, this information is always stored in NVRAM. Use this value in the iSCSI information screen during Oracle Solaris installation so that Oracle Solaris is configured using this ISCSI initiator IQN value.

See Appendix B: DHCP Setup Tips for further information on how to set up DHCP to supply the required network and iSCSI boot information.

Required Configuration Information for Setting Up an iSCSI Connection

In order to set up a connection between a SCSI initiator and target using TCP, the information shown in the following figure is required.

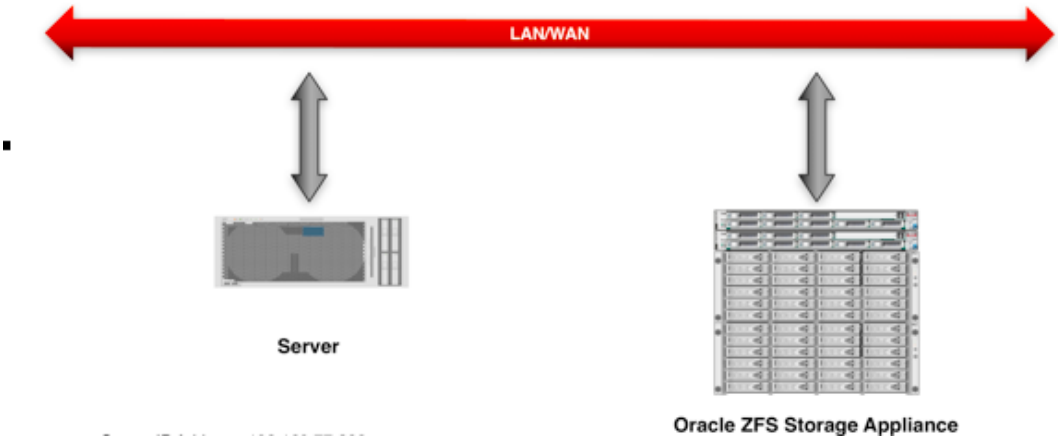

Figure 1. Required information for iSCSI setup

This information is needed both during Oracle Solaris installation and the Oracle Solaris boot process.

Oracle Solaris installation

During Oracle Solaris installation, the installation environment prompts the user to supply the required information for the installation process to discover the LUNs available using iSCSI from the specified Oracle ZFS Storage Appliance IP address. From there it determines the target IQN name and the LUNs available. The Oracle Solaris image to be installed on the LUN is then configured to use the earlier specified iSCSI initiator name. Depending on the network configuration chosen by the user during the installation process, Oracle Solaris is set up to use either a fixed IP address as supplied by the user or DHCP. For SPARC servers, the information needed to boot from the ISCSI LUN is stored in OBP NVRAM.

Oracle Solaris boot

In the Oracle Solaris boot process, it is the actual server boot program that needs access to the previously mentioned information.

Configuring the Oracle ZFS Storage Appliance for iSCSI Boot

The purpose-driven design of the Oracle ZFS Storage Appliance focuses on filesystem solutions while also offering Fibre Channel (FC) and iSCSI block-based solutions. The platform is based on industry-standard x86 architecture, using high-speed memory, Solid State Disk Drives (SSDs) and Serial Attached SCSI (SAS)-based disk storage.

Stability of the platform is ensured by use of Oracle Solaris as a core component, along with Oracle ZFS storage file and block-based services. Oracle Solaris has a long, tested history serving many customers in business-critical environments, utilizing the NFS file sharing protocol, developed by Sun Microsystems, as a key component. The Oracle Solaris operating system has been optimally tuned and configured for the Oracle ZFS Storage Appliance and is preloaded in it as an embedded firmware package, relieving IT administrators from complex OS configuration and upgrade tasks.

When setting up an iSCSI boot environment on an Oracle ZFS Storage Appliance, you must complete the following steps:

- Create a separate project and set required LUN default values.
  - When required, set up a replication policy for the iscsiboot project.
  - When required, set up snapshot policies for the iscsiboot project.
- Create iSCSI initiators and targets and related groups.

  LUNs are exported to the outside world using iSCSI targets. An iSCSI target definition contains one or more network interface ports by which the LUN can be exported. iSCSI initiators and iSCSI targets can be grouped into iSCSI initiator and iSCSI target groups. A LUN is assigned to one iSCSI initiator group and one iSCSI target group. The initiator group is used to limit the LUNs' visibility to only those iSCSI initiators that are members of the initiator group.
- Create a boot LUN for each Oracle Solaris server using the iSCSI boot option and assign it to the just created iSCSI initiator group and iSCSI target group.
- Set up snapshot policies when required.

Appendix C: Configuration Information Handling presents a flowchart showing the overall installation process and its steps, which highlight when particular configuration information is required.

Creating a Project

When setting up an Oracle ZFS Storage Appliance to use an iSCSI boot environment, it is a good practice to create a separate project for that environment so that common configuration information to be used on all LUNs can be established within the project. Figure 2 shows the Create Project pop-up window with the new project name iSCSI_Boot entered.

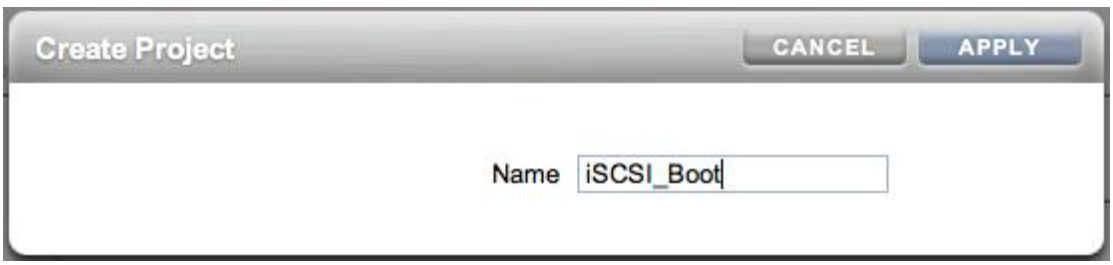

Figure 2. Creating a new project on the Oracle ZFS Storage Appliance

After the project is created, as shown in figure 3, you can set the default configuration properties for the LUNs that will be created within the project.

Figure 3. Setting project properties

Under General, define the defaults to be used when LUNs are created.

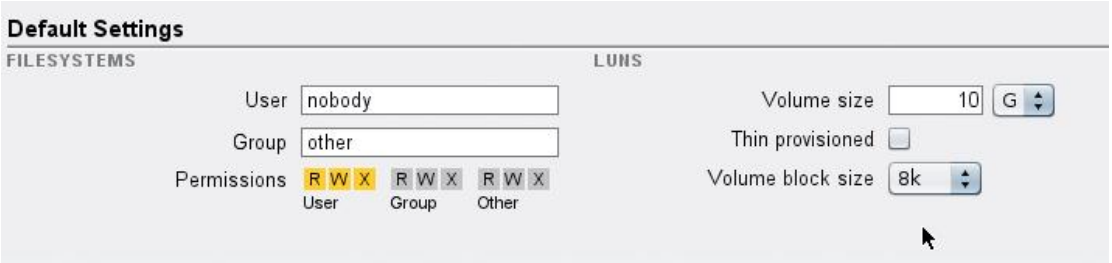

Figure 4. Project LUN defaults

Using a default volume size of 10G provides enough space to hold an Oracle Solaris rpool (root) volume. Increase the size if you would like to add some application and/or user home directory space on the same rpool. Oracle Solaris upgrade procedures need some extra space too, depending on how many install images you want to keep. Oracle Solaris has an automated snapshot feature, TimeSlider, which is useful for restoring data from previous days/months. Using this feature requires extra space for the rpool as well. A default block size of 8K is a good choice for a LUN containing the Oracle Solaris rpool.

For application LUNs, the volume block size depends on the type of I/O. For a random, transactional type of I/O or I/O on lots of small files, it is recommended to keep the volume block size small. For sequential-type I/O with large block transfers, select a larger block size. Once a block size value is set while creating a LUN, the size cannot be changed later.

As this project is only going to be used for iSCSI LUNs, switch off all filesystem type sharing under **Protocols.**

For business continuity reasons, LUNs may need to be replicated elsewhere. Replication can be set up for the project so that all LUNs that are created within the iSCSI_Boot project will automatically be replicated to the secondary Oracle ZFS Storage Appliance.

Creating iSCSI Target and iSCSI Target Groups

Before you can create LUNs you need to define an iSCSI target and an iSCSI target group. An iSCSI LUN target will be exported to one or more network interfaces. In the following example, an IP network multipathing (IPMP) group called `ipmp1` has been created. This group contains two network interfaces, e1000g3 and e1000g4. A LUN using this iSCSI target will be visible on both e1000g3 and e1000g4 network interfaces.

Locate iSCSI target functionality in the Oracle ZFS Storage Appliance BUI by navigating through **Configuration -> SAN ->Targets -> iSCSI Targets.** Use the **+** symbol to create a new iSCSI target.

Select the required options for the new iSCSI target. When CHAP authentication is required, specify the credentials here. The Manual Target IQN option allows you to specify a more logical name. It is recommended to adhere to the `iqn.1986-03.com.sun:` prefix of the IQN name. Names must be unique within your network infrastructure.

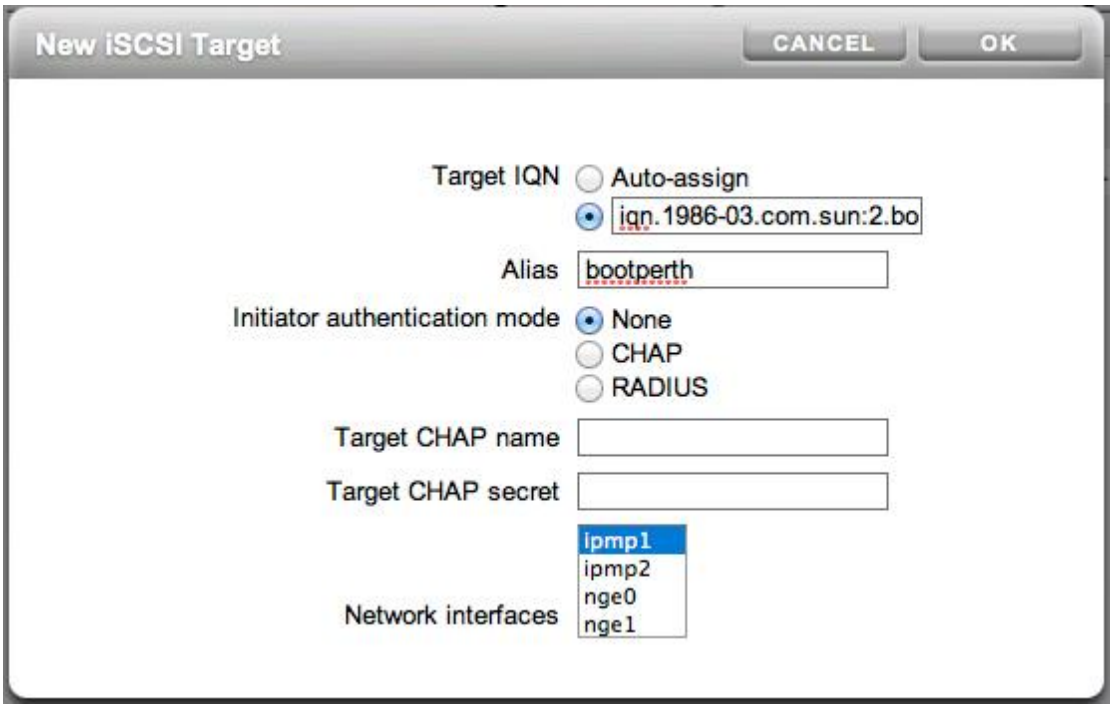

Figure 5. Creating the iSCSI target

The next step is to create a target group. You can simply drag the just-created target to the right under **iSCSI Target Groups**, as seen in the following BUI image. Once the new target group is created, it can be given a more logical name. The example shows two newly created targets and two target groups, one for the SPARC server named perth and one for an x86-type server named x86iscsi.

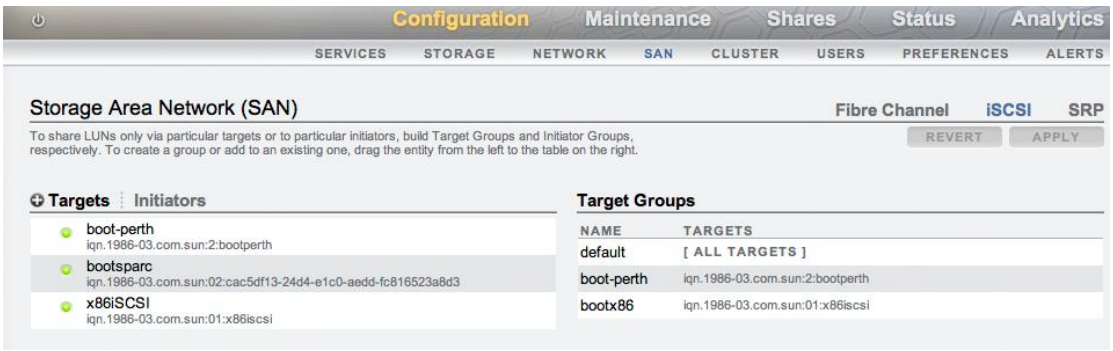

Figure 6. iSCSI targets and target groups

Creating iSCSI Initiators and iSCSI Initiator Groups

To selectively present LUNs to iSCSI initiators requires iSCSI initiator groups. Start by entering the iSCSI IQN name for each server that will be given access to the Oracle ZFS Storage Appliance using the following BUI pop-up window:

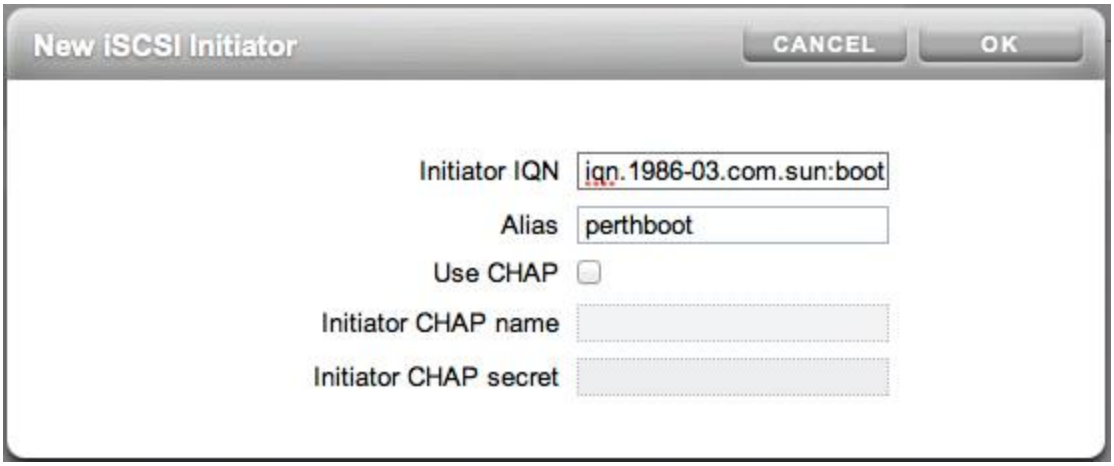

Figure 7. Registering an iSCSI Initiator IQN

Use the same value for the iSCSI initiator IQN that is used in the Oracle Solaris installation process. Follow the syntax `iqn.186-03.com.sun:boot.[`*`NIC MAC HEX value`*`]`. Do not use the CHAP option here as the iSCSI boot environment on the server does not support this option.

To create an iSCSI initiator IQN group, drag the iSCSI initiator from the left column to the right column under **iSCSI Initiator Groups**. The following figure shows iSCSI initiator IQN declarations from three servers and two iSCSI IQN group definitions.

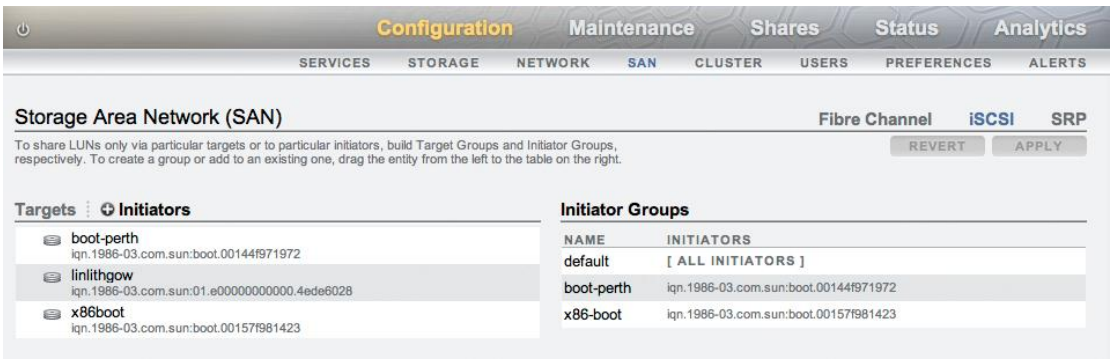

Figure 8. ISCSI Initiator IQN Groups setup

Creating iSCSI Boot LUNs

For each server, you can now create an iSCSI boot LUN in the previously created project. Select the initiator group that contains the iSCSI initiator IQN of the server that is going to use this LUN to boot from, and export the LUN to the target group containing the iSCSI target IQN value the server will use to make the connection to this LUN.

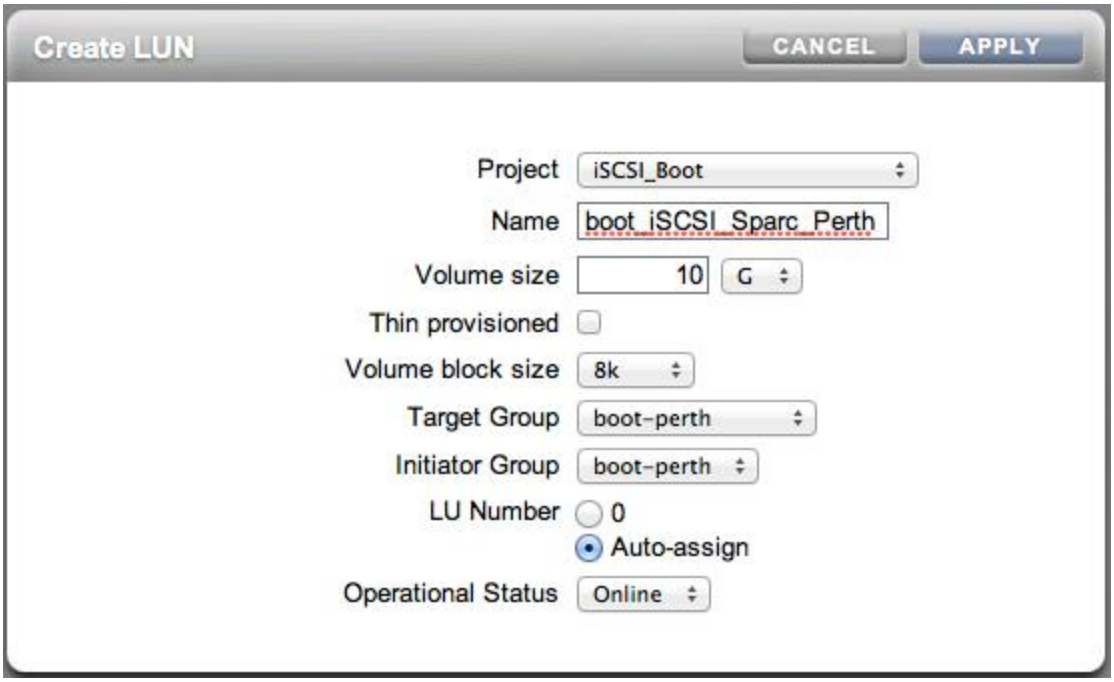

Figure 9. Creating an iSCSI boot LUN

Once the LUN is created, use the Fix LUN Number option in the project's Protocols menu so that the LUN number always matches the LUN number as defined in the server's boot environment.

The following figure shows two LUNs: one used for the SPARC server perth and one for the x86 server x86iSCSI.

Figure 10. Created boot LUNs displayed in Oracle ZFS Storage Appliance BUI

Note the GUID (Globally Unique Identifier) information for each LUN. This value helps identify that the correct LUN is shown in the procedure for installing an Oracle Solaris root.

You are now ready to install Oracle Solaris on the newly configured LUNs.

Installing Oracle Solaris on the Server

The iSCSI installation component to the process for installing Oracle Solaris is identical for both SPARC and x86 hardware platforms.

To perform an Oracle Solaris installation using the iSCSI boot option, the following are needed:

- A CDROM containing the Oracle Solaris **text installation** image or access to a server that provides the Oracle Solaris Automated Installation environment.
- A network connection with, preferably, access to a DHCP service. This DHCP service should be configured to provide a fixed address for the server on which Oracle Solaris will be installed. See Appendix B: DHCP Setup Tips for an example of `dhcpd.conf` information.

  A network connection can be manually configured when a DHCP service is not available. See Appendix A: Manually Configuring a Network Connection.

To start the Oracle Solaris installation process, boot from the CDROM or use the network boot instructions to access the boot image from the Automated Installer (AI) Server as described in the section 'Installing Oracle Solaris 11.1 Systems' from the Oracle Solaris 11.1 Information Library (see Appendix D: References). Follow the **text installation** option in the documentation.

Installing Oracle Solaris on a SPARC Platform

Follow these steps to install Oracle Solaris on a SPARC platform.

1. Shut down the system using:

```
root@solaris:/root# init 0
```

2. Connect to the console using sp.

```
start /SP/console
```

3. If needed, make a note of the current boot device:

```
ok printenv boot-device
boot-device =           /pci@0/pci@0/pci@2/scsi@0/disk@0,0:a
```

4. Boot from the Oracle Solaris Installation image.

   Using a CDROM as the boot medium:

```
Ok boot cdrom: - text-install
```

   Using the network from an Oracle Solaris Automated Installer environment:

```
Ok boot net:dhcp – text-install
```

Installing Oracle Solaris on an x86 Platform

Set the BIOS up to boot from CDROM, or set up a NIC to do a PXE boot when booting using the network from an Oracle Solaris Automated Installer environment.

Installing the iSCSI Component (Platform-independent)

Once the Installer is initiated, the Oracle Solaris installation text procedure brings up the following menu:

```
Welcome to the Oracle Solaris installation menu

        1  Install Oracle Solaris
        2  Install Additional Drivers
        3  Shell
        4  Terminal type (currently xterm)
        5  Reboot
Please enter a number [1]:
```

Use the shell if you need to manually set up the network on the server or if you want to verify access to the configured iSCSI target(s) on the Oracle ZFS Storage Appliance.

Continue the installation process with option 1 from the menu. This will bring up the following screen.

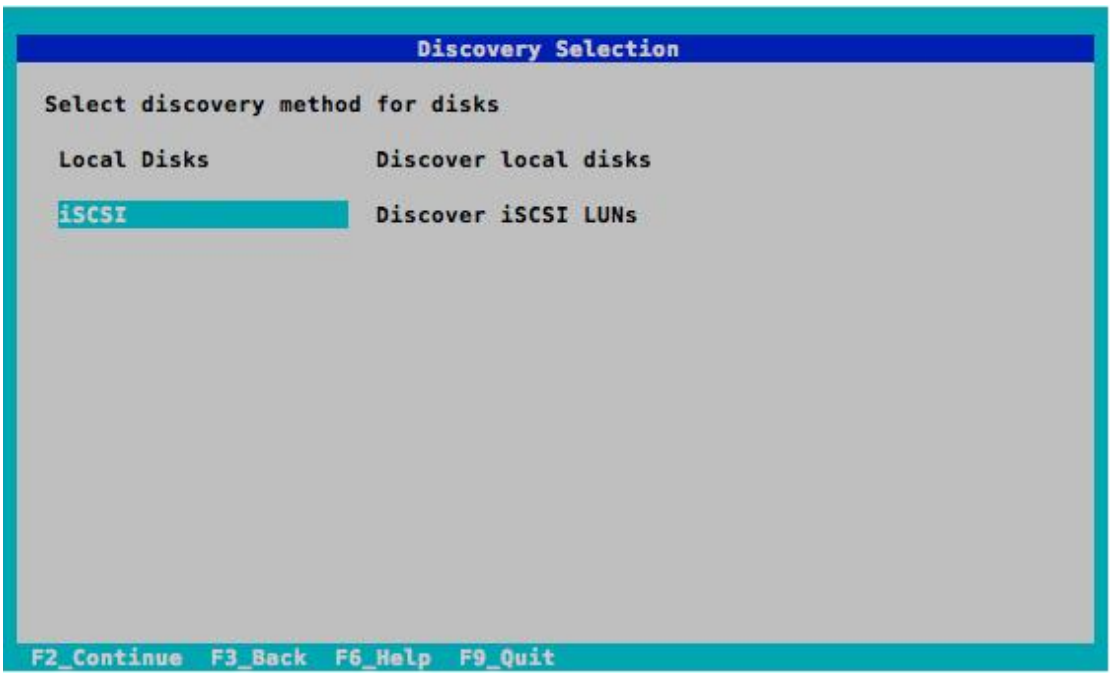

Figure 11. iSCSI LUN option

Using F2 starts the iSCSI installation process. This process will prompt for information to locate the required LUN on which Oracle Solaris is going to be installed, as seen in the following figure.

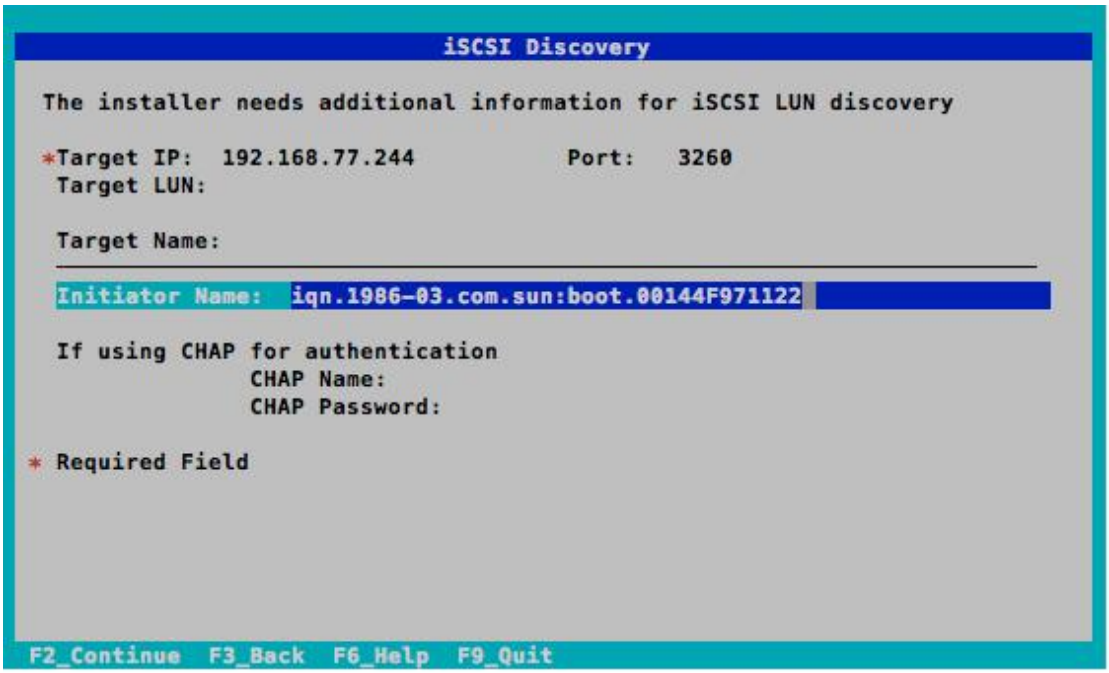

Figure 12. Specifying iSCSI configuration information to locate the target LUN

At this point you need to specify the IP address of the port on the Oracle ZFS Storage Appliance that is used to map the iSCSI target for the boot LUN. Optionally you can specify LUN target access CHAP authentication. Information used here must match the CHAP authentication information set on the Oracle ZFS Storage Appliance for iSCSI target configuration. No CHAP authentication from the Oracle ZFS Storage Appliance to the server can be configured. If more than one LUN is visible from the Oracle ZFS Storage Appliance, you can restrict the number of LUNs being discovered by specifying the LUN number and/or the iSCSI target IQN in the target name field.

Carefully consider what value to use for the Initiator Name. The Oracle Solaris installation process will use this to set up the iSCSI initiator IQN value in the newly installed Oracle Solaris image. When Oracle Solaris boots, it will use this value. As previously mentioned, use the syntax `iqn.1986-03.com.sun:boot.[`*`NIC MAC hex value`*`]`.

After you confirm the iSCSI configuration details, the next window will show all local and iSCSI targets available for an Oracle Solaris installation. The iSCSI LUN(s) will display at the bottom of the list, so when there are multiple local disks shown in the list, keep scrolling down.

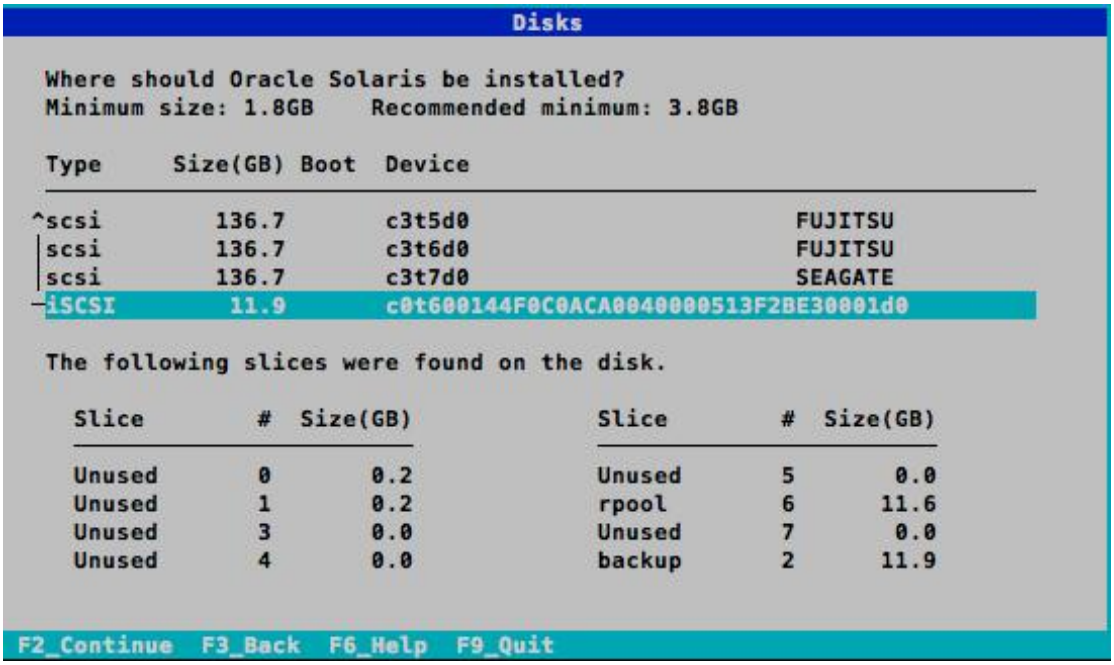

Figure 13. Available disks for Oracle Solaris installation

Note that the device name can be used to verify if the correct LUN has been identified. It should match the GUID string of the LUN as shown in the **SHARES** view of the Oracle ZFS Storage Appliance BUI.

Figure 14. LUN GUID information displayed in the Oracle ZFS Storage Appliance BUI

Confirm the selected disk. The next screen will show the partition information present on the LUN. Always select the whole disk to be used, as seen in the following figure, to avoid any partition alignment issues with the underlying block structure of the LUN.

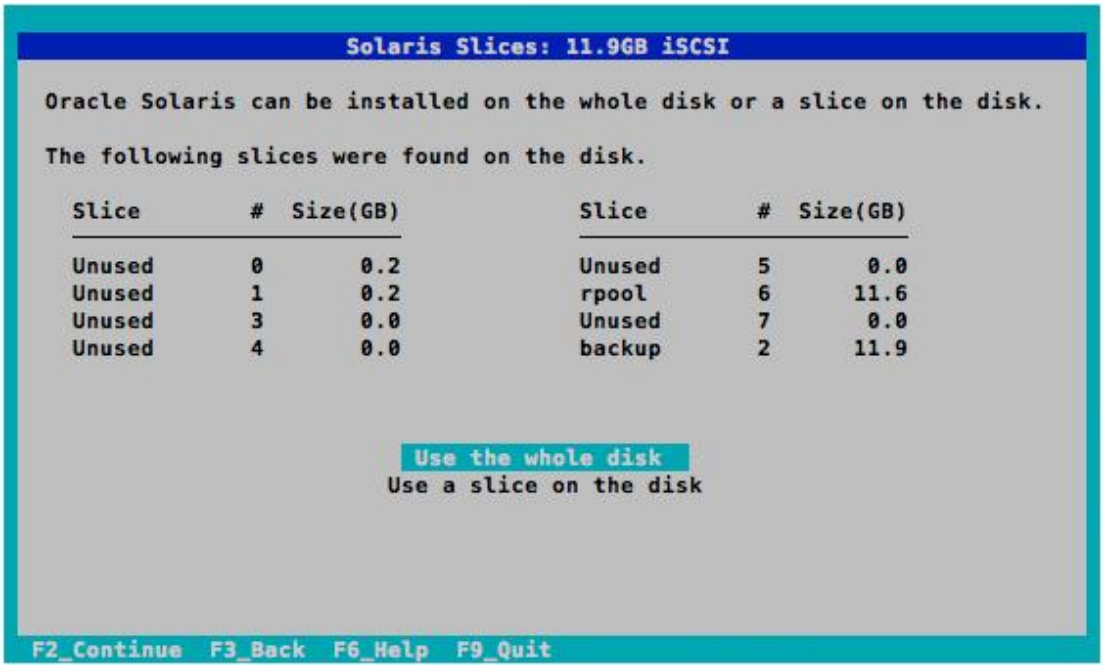

Figure 15. Selected boot disk details

Confirm the selection ‘Use the whole disk’ and press F2 to continue with the rest of the Oracle Solaris installation. Once the installation is in progress, you can use the analytics function on the Oracle ZFS Storage Appliance to verify that the data is transferred to the LUN. Notice that you can see the target group and initiator info on the left side of the traffic graph.

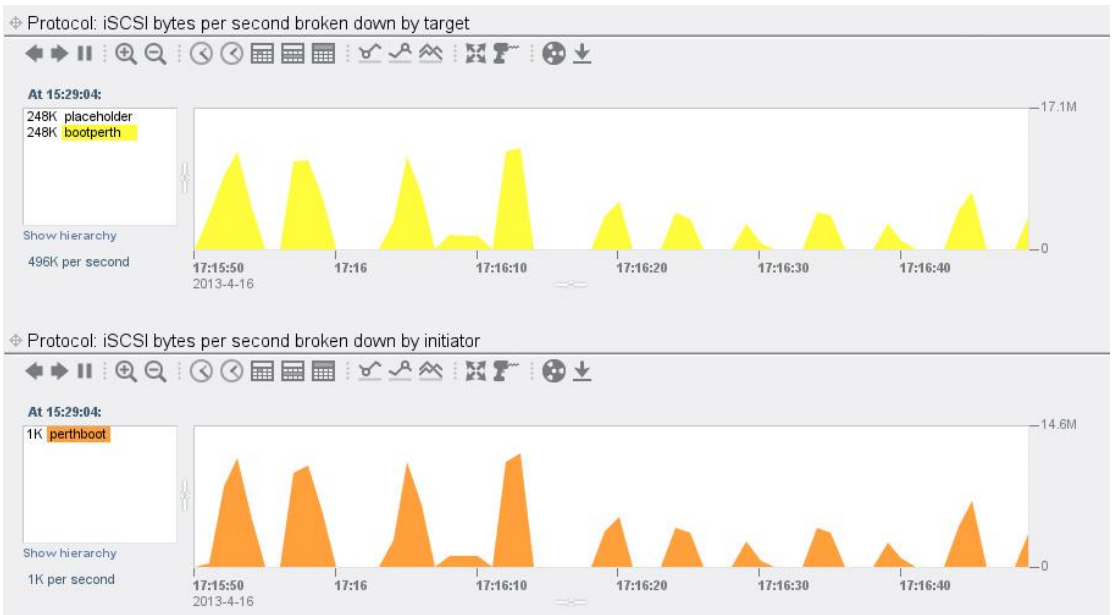

Figure 16. Analytics to show iSCSI install I/O

At the end of the install process, an Oracle Solaris image is installed on the iSCSI LUN. The next step is to configure the server hardware to boot from it. Remember that during the installation process the Oracle Solaris install environment was used to build up an iSCSI connection to the LUN. When rebooting the server, this environment is gone. The hardware must be able to set up a connection to the iSCSI LUN in order to load the

Oracle Solaris boot code. Select 'reboot' at the end of the install process once you are familiar with the procedure as described in the next section.

Configuring the Server for iSCSI Boot

Now that an Oracle Solaris image is installed on the iSCSI LUN, you must next set up the hardware so that it can access its boot loader.

The boot program must be able to initiate an iSCSI connection to the LUN on the Oracle ZFS Storage Appliance to retrieve the Oracle Solaris boot loader. In order to set up an iSCSI session, the boot program must be able to establish its own IP address and iSCSI initiator IQN name. Next, the boot program must determine the access information on the LUN: the Oracle ZFS Storage Appliance IP address, the LUN number, and the IP port through which the Oracle ZFS Storage Appliance accepts iSCSI connection requests.

For the Oracle ZFS Storage Appliance to allow the server to access the LUN, the server must supply its iSCSI initiator IQN as configured earlier in the Oracle ZFS Storage Appliance iSCSI initiator group. In order to reach the Oracle ZFS Storage Appliance, the server needs a correct IP address.

At the end of the Oracle Solaris installation process a reboot is requested. Follow the configuration process for your hardware platform as outlined in the following paragraphs.

Rebooting on the Oracle Solaris SPARC Platform

When rebooting a SPARC server at the end of the Oracle Solaris installation, OBP uses the information from the boot-device and network-boot parameters in the NVRAM. It restarts the server using the OBP `boot net` command.

The OBP network-boot arguments and boot-device variables are updated by the Oracle Solaris installation procedure with the information specified in the installation steps. It will contain information similar to the following:

```
network-boot-arguments=iscsi-target-ip=192.168.77.244,iscsi-target-
name=iqn.1986-03.com.sun:02:boot,host-ip=192.168.77.226,iscsi-
port=3260,iscsi-lun=1,iscsi-partition=a,router-ip=192.168.77.254,subnet-
mask=255.255.255.0.
```

Note that the iSCSI initiator IQN value is missing. Because OBP uses the syntax `iqn.1986-03.com.sun:boot.[`*`NIC MAC hex value`*`]` when the iSCSI initiator IQN value is not defined, the SPARC server can boot from the iSCSI LUN without any further configuration steps.

After the installation is finished, use the Oracle Solaris command `eeprom` to check the configured values.

```
root@solaris:/root# eeprom

~

network-boot-arguments=iscsi-target-ip=192.168.77.244,iscsi-target-
name=iqn.1986-03.com.sun:02:boot,host-ip=192.168.77.226,iscsi-
port=3260,iscsi-lun=1,iscsi-partition=a,router-ip=192.168.77.254,subnet-
mask=255.255.255.0

boot-device=net disk
```

~~

If you would like to use the DHCP service to supply the boot information, you must change a few things in the configuration. Make sure to use the previously mentioned initiator IQN, and use that string when creating iSCSI initiator declarations on the Oracle ZFS Storage Appliance.

In order to boot the server using DHCP, make the following changes:

- Modify the boot-device contents from `net` to `net:dhcp`.
- Set up DHCP to supply the server's network information and the iSCSI boot information. Review Appendix B: DHCP Setup Tips, paying particular attention to the paragraph describing setup when the DHCP is configured to use the Oracle Solaris Automated Installation Service.

Rebooting on the x86 Platform

When you follow the reboot step at the end of the Oracle Solaris installation, intervene in the boot process at the point at which the NIC displays the message "Press <Ctrl-D> to run setup..." as seen in the following figure. Go ahead and press the Control D keys.

Figure 17. Starting NIC iSCSI setup

For this setup routine, you will supply the iSCSI configuration information on the NIC and store it in its NVRAM, then enable the boot option on the NIC.

Select the required port to act as primary boot and press enter to configure the port for iSCSI boot.

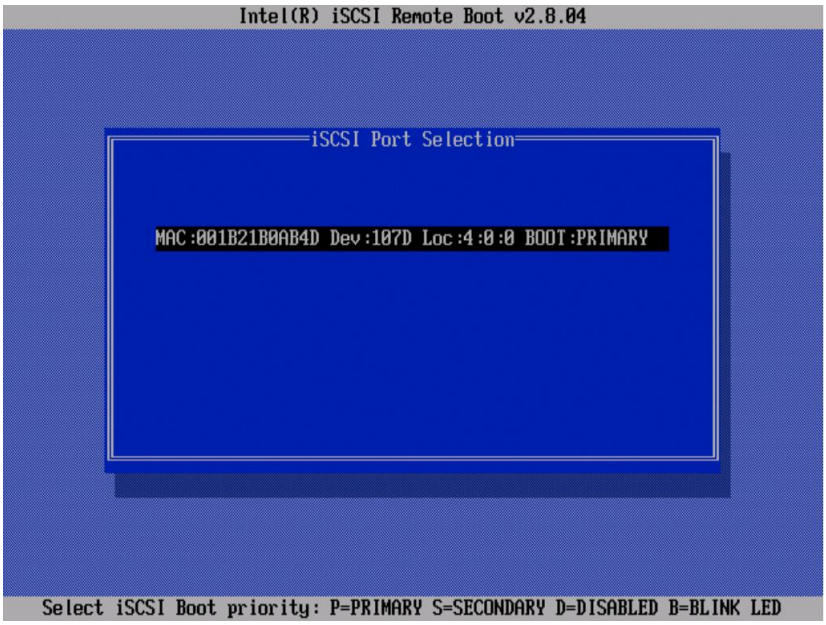

Figure 18. NIC setup; selecting the primary boot port

The next menu contains the options to configure the port, supply optional CHAP login credentials, and store the configuration information to NVRAM of the NIC. Selecting the iSCSI Boot configuration option brings up the following screen.

When using a DHCP service, select the DHCP option for both server IP configuration information and iSCSI target information. Note that you still have to supply the iSCSI initiator IQN name, as the Intel iSCSI remote boot firmware does not currently support retrieving this info from the DHCP service.

When opting not to use a DHCP service, you must supply all configuration information in this setup screen. Select OK to confirm the details.

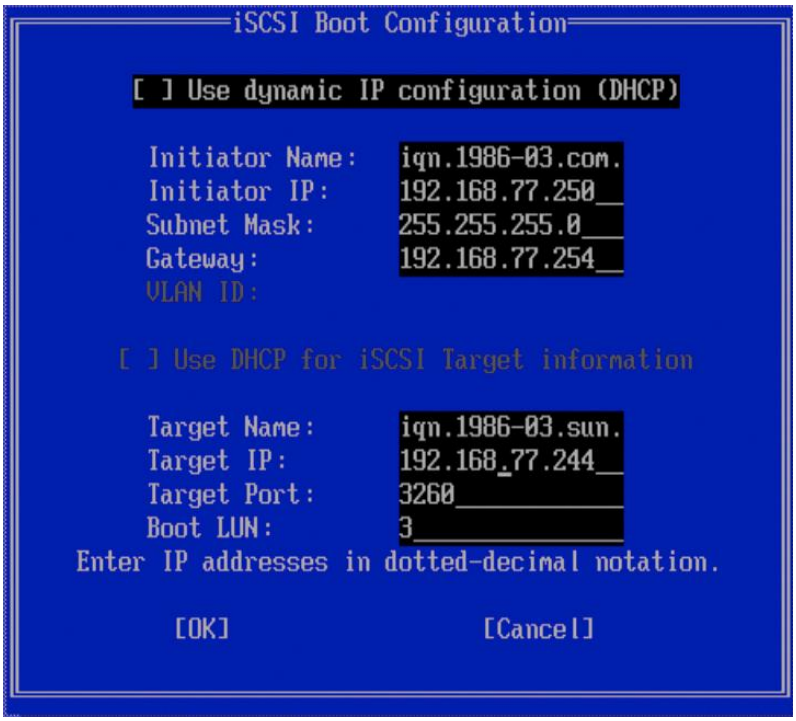

Figure 19. NIC setup; entering iSCSI configuration information

Save the information using “Save changes and Exit” to let the NIC remote boot program connect to the LUN and load the Oracle Solaris boot program.

Figure 20. NIC setup; saving changes

Oracle Solaris Setup

Once Oracle Solaris has come up, log in and verify that the correct iSCSI initiator IQN value has been set using the following command.

```
root@solaris:/root# iscsiadm list initiator-node
Initiator node name: iqn.1986-03.com.sun:boot.00111223344
```

Best Practices for Deploying iSCSI Boot

Consider the following when planning to implement iSCSI boot on the Oracle ZFS Storage Appliance.

Network Topology

Some network configuration options and choices need to be considered when designing a network topology for an iSCSI boot solution.

The Oracle ZFS Storage Appliance contains a number of network ports, which makes it possible to separate administrative, data, and iSCSI network traffic from each other. Doing this ensures that high I/O loads in one area do not impact other services. There are two ways of separating network traffic between different data services: either using virtual networks or allocating physical ports on different subnets.

The Oracle ZFS Storage Appliance supports Link Aggregation Control Protocol (LACP). This way, bandwidth of multiple physical ports can be aggregated into one logical network link. LACP requires use of a network switch that supports this protocol.

Extra availability can be achieved using IP network multipathing (IPMP). IPMP provides redundancy against port/path failure between the Oracle ZFS Storage Appliance and the server. IPMP still uses one logical IP address for access.

It is a good practice to reserve one physical port on each Oracle ZFS Storage Appliance for administrative access.

So what are the best practices for using these options?

The determining factors for deciding on the network topology are the number of servers that are connected to the Oracle ZFS Storage Appliance using iSCSI boot and their expected network load.

If a reasonable load is expected, consider using 10GE ports, or use LACP to aggregate enough 1GE ports to handle the required bandwidth.

If the Oracle ZFS Storage Appliance is used for other types of services (like NFS or CIFS shares), create separate IP subnets for each of them and allocate separate network ports to each subnet.

Using Multipathing

For high availability, consider the use of two separate subnets. Oracle Solaris supports iSCSI LUN multipath access, which is a good alternative to using IPMP. Multipathing provides redundancy against loss of a complete subnet. In contrast, IPMP works within one subnet and thus provides redundancy against loss of network port failures, not the loss of a complete subnet.

LUNs in the Oracle ZFS Storage Appliance can be made visible on multiple network interfaces when creating target definitions. Those targets will be visible to the initiator over the IP interfaces as configured in the target definition.

When using an Oracle ZFS Storage Appliance in a cluster configuration, you must only use networking interfaces active on a single node in a target definition. This restriction also applies to physical network interfaces used in logical LACP and IPMP interfaces.

Make sure that the equivalent ports on the other node are not in use so that when an Oracle ZFS Storage Appliance node fails, all active network ports can fail over onto the new active node.

On the Oracle Solaris server side, make sure the number of iSCSI sessions configured matches the number of path available to the iSCSI target. See the following Oracle Solaris admin docs for more information on how to set up iSCSI multipathing on the server:

"Setting Up iSCSI Multipathed Devices in Oracle Solaris"

http://docs.oracle.com/cd/E23824_01/html/821-1459/fncpc.html#scrolltoc

"How to Enable Multiple iSCSI Sessions for a Target"

http://docs.oracle.com/cd/E19253-01/817-5093/gcawf/

Using DHCP

Using DHCP enables you to centrally manage the IP addresses in your infrastructure. When using DHCP you have the option to use dynamic IP allocation or fixed IP allocation for the Oracle Solaris server that is going to use iSCSI boot.

As described previously in this paper, in some cases the server's IP address is stored in NVRAM of the NIC or OBP (SPARC). In such cases you cannot use dynamic IP allocation. Fixed IP allocation can still be used.

Care should be taken when the Oracle Solaris Automated Installer function is also used by the DHCP service. See Appendix B: DHCP Setup Tips for more information.

Troubleshooting iSCSI Configuration Problems

The following are common potential problems with a new iSCSI configuration, with provided troubleshooting solutions.

No Network Connection Configured During Oracle Solaris Installation

Check if you have a DHCP service available in the subnet to which the server is connected. If so, set up the DHCP server to provide a static IP address for this server by adding the following configuration information in the `dhcpd.conf` file.

The Oracle Solaris installation procedure depends on one network port being active and configured for network access. It uses DHCP to retrieve the needed network configuration information.

If you do not have a DHCP server available on the network your server is connected to, you can configure an IP port for network access manually. You will need the following IP configuration information: client's IP address, netmask, and default router. See Appendix A: Manually Configuring a Network Connection for details for the Oracle Solaris iSCSI installation process.

Need to Verify iSCSI Targets' Access from Initiator

Use the shell in the Oracle Solaris install procedure for the following steps:

1. Add a target address for discovery.

2. Enable static discovery.

3. Verify the initiator-node name.

If there is no initiator node defined, create one. If the name does not match with the name you intended to use, change the name with the `iscsiadm modify initiator-node name` command option.

```
root@solaris:/root# iscsiadm add discovery-address 192.168.77.244

root@solaris:/root# iscsiadm modify discovery –static enable

root@solaris:/root# iscsiadm list initiator-node

Initiator node name: iqn.1986-03.com.sun:boot.00111223344

....

...

root@solaris:/root# iscsiadm list target -S
Target:    iqn.1986-03.com.sun:02:bootx86iscsi
           Alias: bootx86iscsi
           TPGT: 2
           ISID: 4000002a0000
           Connections: 1
           LUN: 3
                 Vendor:   SUN
                 Product:  Sun Storage 7410
                 OS Device Name:
/dev/rdsk/c0t600144F0C0ACA00400005151BFA40002d0s2
root@solaris:/root#
```

The last command should show the target IQN and the LUN upon which you are planning to install the Oracle Solaris OS. If no target information or the wrong target information displays, check the target and initiator group definitions on the Oracle ZFS Storage Appliance.

Before leaving the shell, make sure you restore the iSCSI configuration to the original state by disabling the iSCSI discovery methods and deleting the discovery address(es).

```
root@solaris:/root# iscsiadm modify discovery –static disable
root@solaris:/root# iscsiadm modify discovery –sendtargets disable
root@solaris:/root# iscsiadm remove discovery-address 192.168.77.244
root@solaris:/root# iscsiadm list discovery
Discovery:
        Static: disabled
        Send Targets: disabled
        iSNS: disabled
root@solaris:/root# iscsiadm list discovery-address
root@solaris:/root# exit
```

For SPARC you can also directly check the visibility of the iSCSI target information using the OBP `show-iscsi` command.

```
ok show-iscsi net:iscsi-target-ip=192.160.77.244,host-
ip=192.168.77.226,iscsi-target-name=xx,router-ip=192.168.77.254,subnet-
mask=255.255.255.0
1000 Mbps full duplex Link up
Target: iqn.1986-03.com.sun:02:bootx86iscsi
Lun 3-0-0-0
Disk     SUN     Sun Storage 74101.0     20971520 Blocks, 10 GB
Target: iqn.1986-03.com.sun:2: boot.00144F971122
Lun 0-0-0-0
Disk     SUN     Sun Storage 74101.0     25165824 Blocks, 12 GB
```

Too Many LUNs Showing Up

Check the target and initiator group configuration on the Oracle ZFS Storage Appliance. To restrict visibility of LUNs configured on the Oracle ZFS Storage Appliance, allocate only one iSCSI initiator per initiator group. Secondly, when using a clustered Oracle ZFS Storage Appliance configuration, only allocate IP interfaces to iSCSI targets that are active on the current node; that is, the node the LUN is active on.

Unable to Access Any iSCSI Targets

Follow the previous section on verifying access to iSCSI targets. If this does not help, configure the LUN on the Oracle ZFS Storage Appliance to the default target group. If this resolves the issue, check the initiator IQN value used in the Oracle ZFS Storage Appliance iSCSI initiator and initiator group definitions. Use the syntax `iqn.1986-03.sun.com:boot.[`*`NIC MAC value in HEX`*`]` to specify the initiator IQN, where `1986-03.sun.com` is vendor specific and can be different for your server hardware and/or NIC.

After establishing the connection to the iSCSI target, reconfigure the LUN on the Oracle ZFS Storage Appliance to use the initiator group that contains the proper initiator IQN value.

Appendix A: Manually Configuring a Network Connection

The Oracle Solaris install procedure requires a working network connection for it to connect to the iSCSI LUN to be used to install the Oracle Solaris OS and configure Oracle Solaris for iSCSI boot. The Oracle Solaris boot process requires access to a DHCP server to obtain the needed IP configuration information like server IP address, netmask, default router and iSCSI boot path.

When the installation infrastructure does not have a DHCP service available, it is possible to manually set up a network connection and then start the Oracle Solaris installation procedure.

Make sure to use a static IP address that is permanently allocated to the new server. This IP information needs to be used when configuring the iSCSI boot information. If your server is a SPARC-based machine, the Oracle Solaris installation process will automatically and correctly set up the OBP network-boot-parameters in NVRAM. When using an x86-based server, you need to specify this information in the NIC's BIOS during the server boot process.

Select the shell from the install menu.

```
Hostname: solaris
Welcome to the Oracle Solaris installation menu

        1  Install Oracle Solaris
        2  Install Additional Drivers
        3  Shell
        4  Terminal type (currently xterm)
        5  Reboot
Please enter a number [1]: 3
To return to the main menu, exit the shell
root@solaris:/root #
```

Check availability of network ports and the link status. At least one port should be up and wired to your network interface.

```
root@solaris:/root#  dladm show-phys
LINK              MEDIA                STATE      SPEED  DUPLEX    DEVICE
e1000g1           Ethernet             unknown    0      unknown   e1000g1
e1000g2           Ethernet             unknown    0      unknown   e1000g2
e1000g0           Ethernet             up         1000   full      e1000g0
e1000g3           Ethernet             unknown    0      unknown   e1000g3
root@solaris:/root# dladm show-link
LINK                CLASS     MTU    STATE    OVER
e1000g1             phys      1500   unknown  --
e1000g2             phys      1500   unknown  --
e1000g0             phys      1500   up       --
e1000g3             phys      1500   unknown  --
```

Configure the IP stack on the first network port.

```
root@solaris:/root# ipadm create-addr -T static -a 192.168.77.226/24
e1000g0/v4static
root@solaris:/root# ipadm show-addr
```

```
ADDROBJ           TYPE     STATE        ADDR
e1000g0/v4static  static   ok           192.168.77.226/24
lo0/v4            static   ok           127.0.0.1/8
lo0/v6            static   ok           ::1/128
```

Set a default route.

```
root@solaris:/root# route –p add default 192.168.77.254
```

The Oracle Solaris installation can now be continued by terminating the shell using either the exit command or the keys Control-D. This will bring you back into the installation menu. Use option 1 to start the Oracle Solaris installation process.

Appendix B: DHCP Setup Tips

The following are recommendations for addressing DHCP functionality.

Setting up for configuration with a fixed server IP address

Use a fixed IP address for the machine to be used for iSCSI boot for the network used to access the iSCSI boot LUN.

Not all hardware supports the use of DHCP when setting up an initial connection to the iSCSI target at the start of the boot sequence. In such situations IP information for both the server and the Oracle ZFS Storage Appliance is stored in the boot hardware's NVRAM. Oracle Solaris, once booted, can pick up its IP address from the DHCP service.

Keep things simple by making those IP addresses the same: set up the DHCP service to return a fixed address for the server. Do this by adding the following configuration declaration for each server for which you want to create a fixed IP address. (Comment lines for this paper's example are included in the code.) If a DNS name service is used, make sure the same IP address is used in the DHCP config as is used in the DNS service.

```
# perth , also used for iscsi boot
host aie-t5220-1 {
  hardware ethernet 00:14:4f:97:11:22;
  fixed-address 192.168.77.226;
  # hostname can be used to let solaris pick up the hostname during
installation
  # make sure it matches the info in DNS if used.
  option hostname “perth”;
}
```

Setting up DHCP to provide iSCSI boot disk info and initiator IQN

Once the Oracle Solaris installation has been finished the iSCSI target LUN information can be added in the DHCP server declaration by the form of the variable *`root-path`*. For hardware and/or OS that support the retrieval of the initiator iqn information the dhcp option `iscsi-initiator-iqn` can be added.

```
# Use the following declaration near the top of your dhcpd.conf file
option iscsi-initiator-iqn code 203 = string;

# perth , also used for iscsi boot
host aie-t5220-1 {
  hardware ethernet 00:14:4f:97:11:22;
  fixed-address 192.168.77.226;
  # hostname can be used to let solaris pickup the hostname during installation
  # make sure it matches the info in DNS if used.
  option hostname “perth”;
# syntax ISCSI:<server name or IP>:<protocol>:<port>:<LUN>:<targetname>
# empty fields use defaults port=3260, LUN=0
 option root-path "iscsi:192.168.77.244:::0:iqn.1986-03.com.sun:2:
```

```
boot.00144F971122";
  filename "";
  option iscsi-initiator-iqn = "iqn.1986-03.com.sun:boot.00144F971122";
}
```

Using DHCP for iSCSI Boot in Coexistence with Automated Installer

The Oracle Solaris Automated Installer adds a couple of lines to the `dhcpd.conf` file to support booting an Oracle Solaris installation image using the `boot net:dhcp` option or PXE boot on x86 platforms.

The Automated Installer uses the DHCP config `/etc/inet/dhcpd4.conf` filename to specify the boot image. The Installer uses rules such as the following (ISC-DHCP `switch` statement syntax used).

```
# AI Install DHCP config info

option arch code 93 = unsigned integer 16;


class "SPARC" {
  match if substring ( option vendor-class-identifier,0,5) = "SUNW.";
  switch (substring(option vendor-class-identifier, 5, 5)) {
    case "Ultra":
    case "SPARC":
          filename "http://192.168.77.91:5555/cgi-bin/wanboot-cgi";
          break;
  }
}
# AI X86 boot install via PXEBoot
class "PXEBoot" {
  match if (substring(option vendor-class-identifier, 0, 9) = "PXEClient");
  if option arch = 00:00 {
    filename "default-i386/boot/grub/pxegrub2";
  } else if option arch = 00:07 {
    filename "default-i386/boot/grub/grub2netx64.efi";
  } else {
        filename "";
  }
}
```

Using both the DHCP option filename and root-path for the same server may lead to confusion. This situation arises for SPARC hardware, as the OBP environment does not provide a mechanism for distinguishing between an iSCSI boot or boot using Automated Installer. Both situations require `boot net:dhcp` when you want OBP to use the root-path from the DHCP service. When using the OBP command `boot net`, OBP uses the iSCSI parameters as specified in the NVRAM network-boot-parameters variable.

Comment out the root-path and filename variable declaration in the server IP config block. Once the installation is completed, uncomment the root-path and filename variable in the server IP declaration block to re-establish the boot from iSCSI LUN option.

Remember to (re)start the `dhcpd` daemon process after making changes to the `.conf` file.

When using the Oracle Solaris `dhcpd` daemon use the following commands:

```
#Stop dhcpd daemon ( ipv4 version used in this example)
svcadm disable /network/dhcp/server:ipv4
# Check the syntax of the config file
/usr/lib/inet/dhcpd -t -cf /etc/inet/dhcpd4.conf
# restart dhcpd daemon
svcadm enable /network/dhcp/server:ipv4
# check if dhcpd daemon started
svcs –xv
```

Appendix C: Configuration Information Handling

The following flowchart shows the installation process and at which point what network and iSCSI configuration information is required and stored for the boot environment.

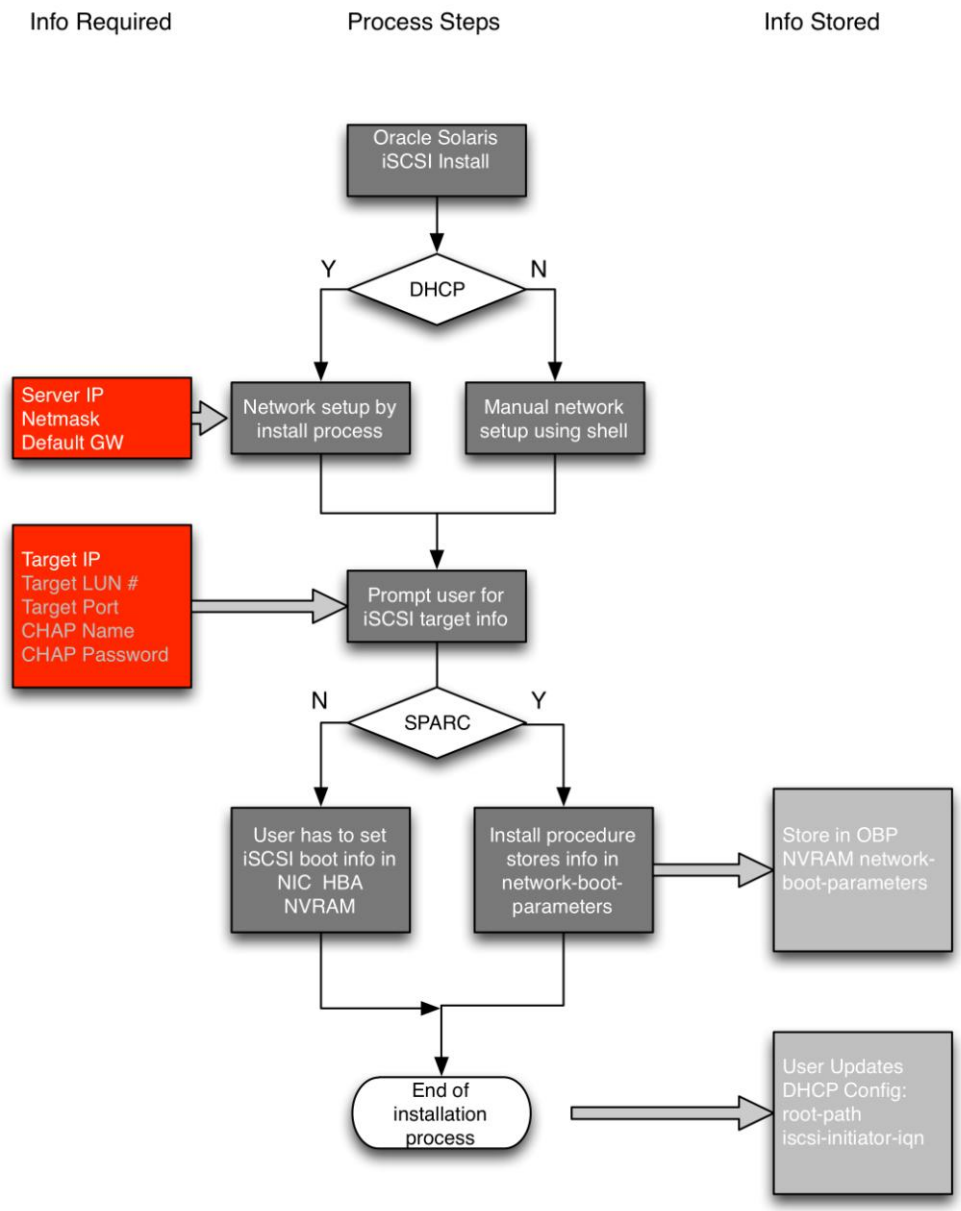

Figure 21. Flowchart for installing Oracle Solaris on an iSCSI LUN

Appendix D: References

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 or screen code may still carry these legacy naming conventions.

Oracle ZFS Storage Appliance Product Information
http://www.oracle.com/us/products/servers-storage/storage/nas/overview/index.html

Oracle ZFS Storage Appliance White Papers and Subject-Specific Resources
http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/index.html

Oracle ZFS Storage Appliance technical documentation
http://docs.oracle.com/cd/E26765_01/index.html

The *Sun ZFS Storage Appliance Administration Guide* is also available through the Oracle ZFS Storage Appliance help context.
The Help function in Oracle ZFS Storage Appliance can be accessed through the browser user interface.

"Simplified Installation and Cloud Provisioning with Oracle Solaris 11" information page
http://www.oracle.com/technetwork/server-storage/solaris11/technologies/modernizedinstaller-461041.html

Oracle Solaris installation options from Oracle Solaris 11.1 Information Library
http://docs.oracle.com/cd/E26502_01/html/E28980/usemedia.html

“How to Install the Oracle Solaris Release on an iSCSI LUN”, Oracle Solaris Device Administration section 821-1459
http://docs.oracle.com/cd/E23824_01/pdf/821-1459.pdf

Implementing iSCSI Boot with Oracle Solaris on the Oracle ZFS Storage Appliance
January 2014; Version 1.0
Author: 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

Hardware and Software, Engineered to Work Together