

How to Implement Oracle Solaris Zones Using the Sun ZFS Storage Appliance

May 2012

By Andrew Ness

This article describes how to implement Oracle Solaris Zones to create virtual Oracle Solaris instances using the Sun ZFS Storage Appliance.

The Oracle Solaris Zones feature offers an isolation environment for Oracle Solaris 11, providing secure, flexible, scalable and lightweight virtualized operating system services which appear as distinct Oracle Solaris instances.

In order to provide this virtualized environment, back-end storage space is required to store not only the Solaris personalization files and zone framework but also any application binaries and data. The Sun ZFS Storage Appliance provides a highly flexible platform that manages this required storage, in a flexible environment that concurrently caters to a range of network infrastructures.

Contents

[Overview](#)

[Choosing iSCSI or Fibre Channel Protocol](#)

[Creating Block-Protocol LUNs on the Sun ZFS Storage Appliance](#)

[Creating the Zone Path LUN on the Sun ZFS Storage Appliance](#)

[Creating a Pool](#)

[Creating a Zone](#)

[Installing the Zone Package Files](#)

[Booting and Logging In to the Zone](#)

[Adding Additional Storage Resources to the Zone](#)

[Using a Block Volume for a Non-Global Zone ZFS Dataset](#)

[Allocating a Non-Global Zone Block Device](#)

[Conclusion](#)

[References](#)

Overview

The Oracle Solaris Zones feature presents a way to segregate applications in the same Oracle Solaris operating system, in an environment that provides allocation control for resources such as CPU, memory, networking bandwidth and storage.

One of the requirements for the storage provisioning for Oracle Solaris Zones is that storage must be local – meaning that NFS-mounted file systems cannot be used. Storage on which the zones are installed must be block storage. By providing both file shares and block-level LUNs, the Sun ZFS Storage Appliance is highly suitable as the back-end storage device.

This article describes the steps required to create an Oracle Solaris Zone on a block LUN presented by the Sun ZFS Storage Appliance, and how to supplement the storage allocation by either presenting additional ZFS storage pools or raw volumes to the zone.

The Sun ZFS Storage Appliance provides a flexible environment to provide storage resources in file or block formats to multiple servers concurrently, so you can rapidly deploy and efficiently share storage resources.

Choosing iSCSI or Fibre Channel Protocol

The choice of which block-level protocol and infrastructure to use depends upon local policy and any existing connections between the Sun ZFS Storage Appliance and the Solaris server.

Fibre Channel SAN deployment requires installation, configuration and management of Fibre Channel switches, cabling, and host bus adapters (HBAs) but offers potentially more consistent performance and segregation between standard network traffic and data access traffic. While Fibre Channel protocol provides deterministic performance at the network level with bandwidth but adds additional administration overhead.

iSCSI deployment can be undertaken with minimal additional configuration other than the usual networking connectivity between storage and server. However, this comes at a performance price, as the standard network and data access traffic contend for resources on the same network. In practice, it may be advisable to provide a second network for routing only iSCSI data, keeping the user and data traffic separate. iSCSI can be deployed over standard Ethernet interfaces, and specialized TCP/IP offload engine (T.O.E.) cards are available to reduce the load on the processors.

The choice of which protocol and, therefore, which network infrastructure to use is not within the scope of this document.

Creating Block-Protocol LUNs on the Sun ZFS Storage Appliance

Once you have selected which protocol to use, you will create the LUNs on the Sun ZFS Storage Appliance and configure them on the Solaris host.

A tutorial on creating iSCSI LUNs for use in a Solaris environment is available at:

<http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-110-iscsi-luns-solaris-519668.html>

An equivalent Fibre Channel LUN tutorial is available at:

<http://www.oracle.com/technetwork/articles/servers-storage-admin/o12-012-fc-lun-solaris-1534203.html>

Whichever block protocol you use, create the LUN with a large block size to allow the ZFS cache to operate efficiently. Ultimately, the block size should be matched to the file system work profile, which will differ according to which applications will reside within the zones. Select a block size of 128KB for the zone path.

Creating the Zone Path LUN on the Sun ZFS Storage Appliance

Once you have configured the LUN and presented it to the Oracle Solaris environment, the LUN will be represented by device nodes in `/dev/dsk` and in `/dev/rdsk`.

Following the example from the iSCSI tutorial mentioned in the preceding section, the LUN will be accessible through `/dev/dsk/c0t600144F08F1579D600004E1C2E250001d0s2` with the corresponding `/dev/rdsk` entry.

Creating a Pool

Using the examples in the iSCSI and FC LUN tutorials, the next steps show you how to create a ZFS pool using the LUNs as the underlying device.

1. First, connect to the Oracle Solaris server that will host the zone to a CLI session. Log in as a valid user and assume the 'root' role.

```
login: an
Password: xxxxxxxxxxxxxx
Last login: DDD MMM DD HH:MM:SS from sgdproxy
an@solaris$ su
Password: xxxxxxxxxxxxxxxxxxxx
```

2. Next, create a pool from which ZFS file systems will be provisioned.

```
root@solaris# zpool create zonepool \
c0t600144F08F1579D600004E1C2E250001d0
```

Now that the pool has been created, you can create the zone configuration and perform the zone installation and first boot.

Creating a Zone

Still logged in as root, you create the zone configuration using the `zonecfg (1m)` command.

The following code sets the zone to automatically boot when the host machine is rebooted by setting the `autoboot` variable to `true`. Since this zone will provide a network service, you can restrict the IP address which it can use by setting the `allowed-address` variable of the network, as seen in the example.

```
root@solaris# zonecfg -z apachezone
apachezone: No such zone configured
Use 'create' to begin configuring a new zone.
zonecfg:apachezone> create
create: Using system default template 'SYSdefault'
zonecfg:apachezone> set zonepath=/zonepool/apachezone
zonecfg:apachezone> set autoboot=true
zonecfg:apachezone> info
zonename: apachezone
zonepath: /zonepool/apachezone
brand: solaris
autoboot: true
bootargs:
file-mac-profile:
pool:
limitpriv:
scheduling-class:
ip-type: exclusive
hostid:
fs-allowed:
anet:
    linkname: net0
    lower-link: auto
    allowed-address not specified
    configure-allowed-address: true
```

```

defrouter not specified
allowed-dhcp-cids not specified
link-protection: mac-nospoof
mac-address: random
mac-prefix not specified
mac-slot not specified
vlan-id not specified
priority not specified
rxrings not specified
txrings not specified
mtu not specified
maxbw not specified
rxfanout not specified
zonecfg:apachezone> select anet linkname=net0
zonecfg:apachezone:anet> set allowed-address=192.168.1.79
zonecfg:apachezone:anet> end
zonecfg:apachezone> verify
zonecfg:apachezone> commit
zonecfg:apachezone> exit

```

Installing the Zone Package Files

Now that the zone has been configured, you must install the necessary packages and configuration, as shown in the following CLI session. Note that as of the Solaris 11 release, the zone package files come from the Oracle Solaris 11 Image Package System (IPS) repository. The installation time depends on the network connectivity between the host server and the IPS repository.

```

root@solaris# zoneadm -z apachezone install
A ZFS file system has been created for this zone.
Progress being logged to
/var/log/zones/zoneadm.20120214T182123Z.apachezone.install
Image: Preparing at /zonepool/apachezone/root.

Install Log: /system/volatile/install.23931/install_log
AI Manifest: /tmp/manifest.xml.v2aiUU
SC Profile: /usr/share/auto_install/sc_profiles/enable_sci.xml
Zonename: apachezone
Installation: Starting ...

Creating IPS image
Installing packages from:
solaris
origin:
http://pkg.oracle.com/solaris/support/
DOWNLOAD          PKGS          FILES          XFER (MB)
Completed          167/167      32064/32064    175.8/175.8

PHASE              ACTIONS
Install Phase      44313/44313

PHASE              ITEMS
Package State Update Phase  167/167
Image State Update Phase    2/2
Installation: Succeeded

```

Note: Man pages can be obtained by installing
pkg:/system/manual

Done: Installation completed in 242.626 seconds.

Next Steps: Boot the zone, then log into the zone console
(zlogin -C) to complete the configuration process.

Log saved in non-global zone as
/zonepool/apachezone/root/var/log/zones/
zoneadm.20120214T182123Z.apachezone.install

```
root@solaris# df -hl | grep zonepool
rpool/zonepool          228G   32K   179G    1%   /zonepool
rpool/zonepool/apachezone 228G 32K   179G    1%
/zonepool/apachezone
rpool/zonepool/apachezone/rpool/ROOT/solaris 228G   310M
179G    1%   /zonepool/apachezone/root
rpool/zonepool/apachezone/rpool/ROOT/solaris/var 228G   26M
179G    1%   /zonepool/apachezone/root/var
```

Booting and Logging In to the Zone

1. Once the zone packages have been successfully installed from the IPS repository, you must boot and log in to the zone to complete the configuration. When connected, you simply follow and respond to the provided prompts.

For details on this procedure, refer to a Solaris 11 zones introduction at:
<http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-092-s11-zones-intro-524494.html>

```
root@solaris# zoneadm -z apachezone boot
root@solaris# zlogin -C apachezone
```

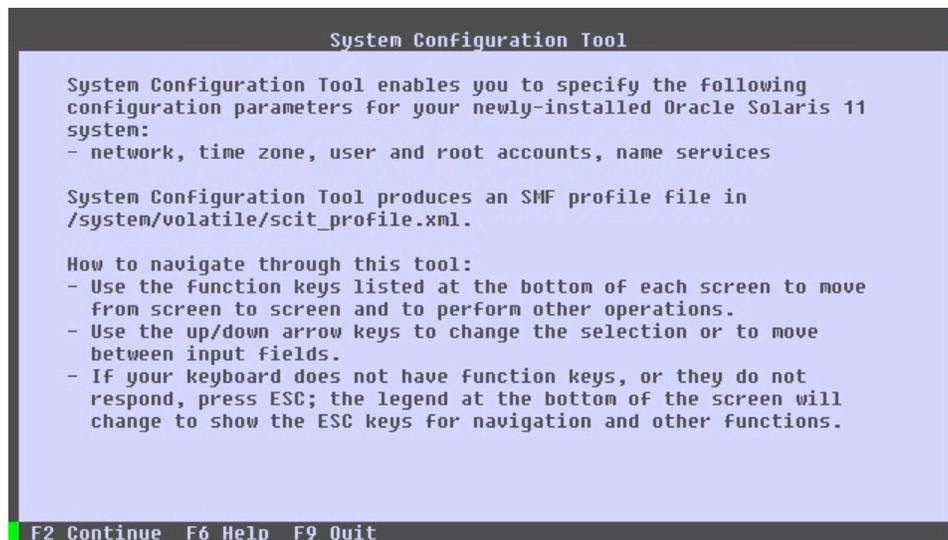


Figure 1. Connecting to the zone console and the system configuration tool

First, provide a name for your zone. In the example, the zone is called apachezone.

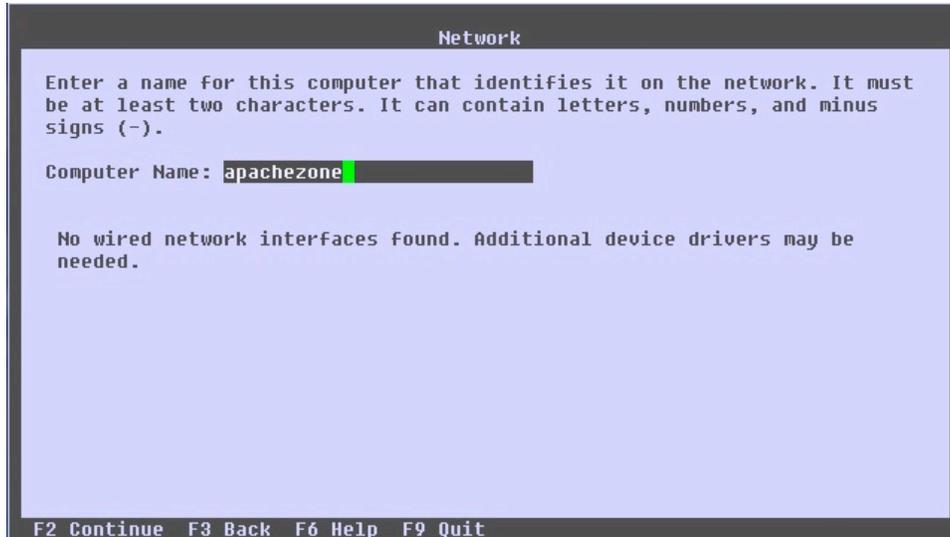


Figure 2. Giving the zone a name

Set up your time zone by selecting your region, your country, and the time zone within that country. The following screen shows `Europe` highlighted as the selected region.



Figure 3. Assigning the time zone by region

The next example screen shows Britain (UK) as the selected country.

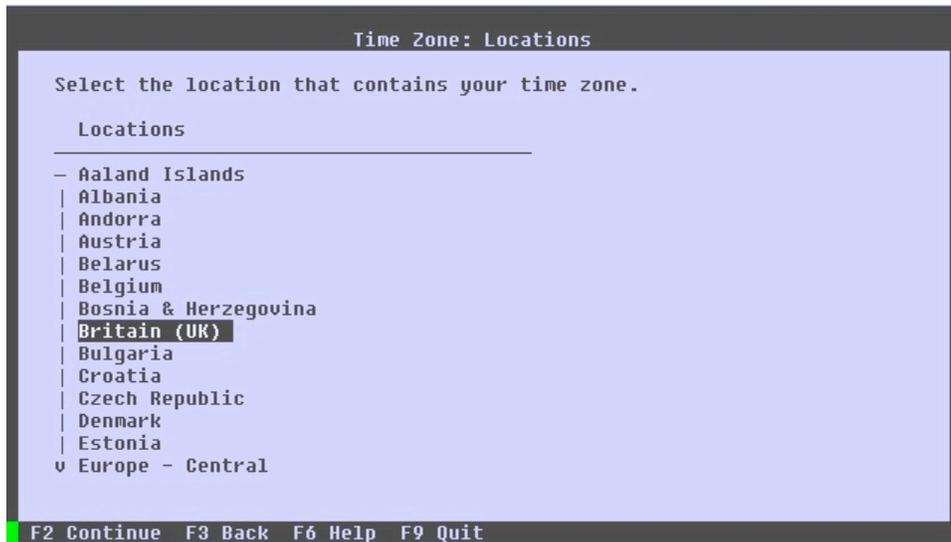


Figure 4. Assigning the time zone by country

The next screen shows the Time Zone selection screen, with GB chosen.



Figure 5. Assigning the time zone

The screen session will prompt you for a root password.

```
Users
Define a root password for the system and user account for yourself.

System Root Password
Root password: *****
Confirm password: *****

Create a user account
Your real name: Joe User
Username: joe
User password: *****
Confirm password: *****

F2 Continue F3 Back F6 Help F9 Quit
```

Figure 6. Assigning the root password

After reviewing your settings, follow the screen instructions to make any needed changes before proceeding.

```
System Configuration Summary
Review the settings below before continuing. Go back (F3) to make changes.

Time Zone: GB

Language: *The following can be changed when logging in.
Default language: C/POSIX
Terminal type: vt100

Users:
Username: joe

Network:
Computer name: apachezone
Manual Configuration: None
IP Address: None
Netmask: 255.255.255.0

F2 Apply F3 Back F6 Help F9 Quit
```

Figure 7. Summary of system configuration

The following screen shows the results of applying the system configuration.

```
Exiting System Configuration Tool. Log is available at:
/var/tmp/install/sysconfig.log
Hostname: apachezone
May  8 14:50:39 apachezone sendmail[5642]: My unqualified host name (apachezone)
unknown; sleeping for retry
May  8 14:50:39 apachezone sendmail[5645]: My unqualified host name (apachezone)
unknown; sleeping for retry

apachezone console login: joe
Password:
Oracle Corporation      SunOS 5.11      11.0      November 2011
joe@apachezone:~$ su
Password:
May  8 14:51:10 apachezone su: 'su root' succeeded for joe on /dev/console
root@apachezone:~# format
Searching for disks...done
No disks found?

root@apachezone:~# █
```

Figure 8. Screen listing the configured zone

2. The environment is now configured to the level where you can install any application or copy any data over to the zone.

Adding Additional Storage Resources to the Zone

Once the zone has been successfully set up, you may want to allocate some dedicated storage as a block device or devices for exclusive use by the zone – for instance, if additional storage is required for application data residing on a raw partition.

Alternatively, if further file system space is required, you can allocate this by adding a ZFS dataset to the non-global zone.

IMPORTANT: Do not build ZFS file systems on LUNs presented from clustered Sun ZFS Storage Appliances. For this setup, you should create UFS file systems.

Using a Block Volume for a Non-Global Zone ZFS Dataset

To provide a ZFS dataset distinct from the root zone dataset, first create the block volume as described in either the Solaris iSCSI or Fibre Channel tutorials mentioned earlier.

Once you have created and assigned the block volume, use it to form a ZFS pool similarly to the previously created zone root pool.

1. As the root role, prepare the LUN for use by ZFS. In order to find the Unique ID of the block LUN, use the `format(1m)` command.

```
root@solaris# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
    0. c0t600144F098D3888900004F3E742D0002d0 <SUN-Sun Storage NAS-
1.0-8.00GB>
       /scsi_vhci/disk@g600144f098d3888900004f3e742d0002
```

```

1. c0t600144F098D3888900004F3E73310001d0 <SUN-Sun Storage NAS-
1.0-24.00GB>
   /scsi_vhci/disk@g600144f098d3888900004f3e73310001
2. c1t0d0 <ATA-ST3500630NS-G-465.76GB>
   /pci@0,0/pci108e,534a@8/disk@0,0
3. c1t1d0 <ATA-ST3500630NS-G-465.76GB>
   /pci@0,0/pci108e,534a@8/disk@1,0
4. c3t0d0 <ATA-SEAGATE ST32500N-.AZK cyl 30397 alt 2 hd 255 sec
63>
   /pci@0,0/pci108e,534a@7/disk@0,0
5. c3t1d0 <ATA-ST3250620AS-E cyl 30397 alt 2 hd 255 sec 63>
   /pci@0,0/pci108e,534a@7/disk@1,0
Specify disk (enter its number): ^D

```

- Next, create a pool from the prepared device. The name of the ZFS file system (apachezonedata in this example) will then appear within the Oracle Solaris Zone as a zpool, as shown in this and the next step.

```

root@solaris# zpool create datapool \
c0t600144F098D3888900004F3E742D0002d0
root@solaris# zfs create datapool/apachezonedata
root@solaris# zoneadm -z apachezone halt
root@solaris# zonecfg -z apachezone
zonecfg:apachezone> add dataset
zonecfg:apachezone:dataset> set name=datapool/apachezonedata
zonecfg:apachezone:dataset> end
zonecfg:apachezone> exit
root@solaris# zonecfg -z apachezone verify

```

- Now boot and log in to the zone.

```

root@solaris# zoneadm -z apachezone boot
root@solaris# zlogin -C apachezone
[Connected to zone 'apachezone' console]

apachezone console login: root
Password:
Feb 17 16:42:55 apachezone login: ROOT LOGIN /dev/console
Last login: Fri Feb 17 16:19:12 on console
Oracle Corporation      SunOS 5.11      11.0      December 2011
root@apachezone:~# zpool list
NAME                SIZE  ALLOC  FREE  CAP  DEDUP  HEALTH  ALTROOT
apachezonedata      7.94G  124K   7.94G  0%   1.00x  ONLINE  -
rpool                23.9G  388M   23.5G  1%   1.00x  ONLINE  -

```

- Next, create a ZFS file system to hold the Apache data. In this example, it is mounted as /apachedata.

```

root@apachezone:~# zfs list
NAME                USED  AVAIL  REFER  MOUNTPOINT
apachezonedata      31K   7.81G   31K    /apachezonedata
rpool               387M   23.1G   31K    /rpool
rpool/ROOT          387M   23.1G   31K    legacy
rpool/ROOT/solaris  387M   23.1G   355M    /
rpool/ROOT/solaris/var 26.9M  23.1G   25.9M   /var
rpool/export        63K   23.1G   32K    /export
rpool/export/home   31K   23.1G   31K    /export/home
root@apachezone:~# zfs create -o mountpoint=/apachedata \

```

```

                                apachezonedata/apachedata
root@apachezone:~# zfs list | grep -i zone
apachezonedata                62K  7.81G   31K  /apachezonedata
apachezonedata/apachedata     31K  7.81G   31K  /apachedata

```

5. If necessary, you can modify ZFS attributes from the non-global zone.

```

root@apachezone:~# zfs set compression=gzip \
                    apachezonedata/apachedata
root@apachezone:~# zfs get compression apachezonedata/apachedata

NAME                                PROPERTY      VALUE          SOURCE
apachezonedata/apachedata          compression   gzip           local

```

Allocating a Non-Global Zone Block Device

To allocate a block device to a non-global zone, first configure the block device in the global zone (described in either the Solaris iSCSI or Fibre Channel tutorials mentioned earlier). Then, after halting the non-global zone to which you will allocate the block device, use the `zonecfg` command to assign the block device.

The volume to be added is created in the same way as previously described (see the section "Creating Block-Protocol LUNs on the Sun ZFS Storage Appliance") and assigned to the global zone for the purposes of iSCSI and FC mapping.

Note: If you are using a file system, it is more efficient to allocate file system objects to the zone instead of the raw partition.

1. Ensure that the zone is halted.

```

root@apachezone:~# halt
Feb 17 15:51:37 apachezone halt: initiated by root on /dev/console

[NOTICE: Zone halted]
~.
[Connection to zone 'apachezone' console closed]

```

2. From the hosting server, prepare the LUN for block device use.

```

root@solaris:~# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
   0. c0t600144F098D3888900004F3E742D0002d0 <SUN-Sun Storage
NAS-1.0 cyl 4094 alt 2 hd 128 sec 32>
      /scsi_vhci/disk@g600144f098d3888900004f3e742d0002
   1. c0t600144F098D3888900004F3E73310001d0 <SUN-Sun Storage
NAS-1.0-24.00GB>
      /scsi_vhci/disk@g600144f098d3888900004f3e73310001
   2. clt0d0 <ATA-ST3500630NS-G-465.76GB>
      /pci@0,0/pci108e,534a@8/disk@0,0
   3. clt1d0 <ATA-ST3500630NS-G-465.76GB>
      /pci@0,0/pci108e,534a@8/disk@1,0
   4. c3t0d0 <ATA-SEAGATE ST32500N-.AZK cyl 30397 alt 2 hd 255
sec 63>
      /pci@0,0/pci108e,534a@7/disk@0,0

```

```

5. c3t1d0 <ATA-ST3250620AS-E cyl 30397 alt 2 hd 255 sec 63>
   /pci@0,0/pci108e,534a@7/disk@1,0
Specify disk (enter its number): 0
selecting c0t600144F098D3888900004F3E742D0002d0
[disk formatted]
No Solaris fdisk partition found.

```

FORMAT MENU:

```

disk          - select a disk
type          - select (define) a disk type
partition     - select (define) a partition table
current       - describe the current disk
format        - format and analyze the disk
fdisk         - run the fdisk program
repair        - repair a defective sector
label         - write label to the disk
analyze       - surface analysis
defect        - defect list management
backup        - search for backup labels
verify        - read and display labels
save          - save new disk/partition definitions
inquiry       - show disk ID
volname       - set 8-character volume name
!<cmd>       - execute <cmd>, then return
quit

```

```
format> fdisk
```

No fdisk table exists. The default partition for the disk is:

```
a 100% "SOLARIS System" partition
```

Type "y" to accept the default partition, otherwise type "n" to edit the partition table.

```
y
```

```
format> ^D
```

3. Add the appropriate device nodes to the zone configuration.

```

root@solaris# zonecfg -z apachezone
zonecfg:apachezone> add device
zonecfg:apachezone:device> set \
match=/dev/dsk/c0t600144F098D3888900004F3E742D0002d0s*
zonecfg:apachezone:device> set allow-partition=true
zonecfg:apachezone:device> set allow-raw-io=true
zonecfg:apachezone:device> end
zonecfg:apachezone> add device
zonecfg:apachezone:device> set \
match=/dev/rdisk/c0t600144F098D3888900004F3E742D0002d0s*
zonecfg:apachezone:device> set allow-partition=true
zonecfg:apachezone:device> set allow-raw-io=true
zonecfg:apachezone:device> end
zonecfg:apachezone> commit
zonecfg:apachezone> exit

```

Note that you must provide both /dev/dsk and /dev/rdisk entries.

4. Next, boot the zone again.

```

root@solaris:~# zoneadm -z apachezone boot
root@solaris:~# zlogin -C apachezone
[Connected to zone 'apachezone' console]
Hostname: apachezone
apachezone console login: root
Password:
Feb 17 16:00:55 apachezone login: ROOT LOGIN /dev/console
Last login: Fri Feb 17 15:48:30 on console
Oracle Corporation      SunOS 5.11      11.0      December 2011

```

5. Now ensure that you have access to the LUN.

```

root@apachezone:~# format
Searching for disks...done

```

```

AVAILABLE DISK SELECTIONS:
    0. c0t600144F098D3888900004F3E742D0002d0 <SUN-Sun Storage
NAS-1.0 cyl 4093 alt 2 hd 128 sec 32>
        sd16 at scsi_vhci0 slave 0
Specify disk (enter its number): ^D
root@apachezone:~#

```

Now you can use the LUN for block access as needed.

Conclusion

The Sun ZFS Storage Appliance provides a highly flexible platform on which to implement the advanced Oracle Solaris features of Zones. The appliance provides not only the required storage for the zone path but also any additional data and application space required in structured ZFS or block access formats.

References

For more information, visit the Web resources listed in Table 1.

Table 1. Web resources for further information	
Web Resource Description	Web Resource URL
How to access iSCSI LUNs in an Oracle Solaris Environment	http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-110-iscsi-luns-solaris-519668.html
How to Make Fibre Channel Storage Available to Oracle Solaris	http://www.oracle.com/technetwork/articles/servers-storage-admin/o12-012-fc-lun-solaris-1534203.html
How to Get Started Creating Oracle Solaris Zones in Oracle Solaris 11	http://www.oracle.com/technetwork/articles/servers-storage-admin/o11-092-s11-zones-intro-524494.html
Oracle Solaris	www.oracle.com/solaris
Sun ZFS Storage Appliances	www.oracle.com/us/products/servers-storage/storage/nas/overview/index.html
Download Oracle Solaris 11	http://www.oracle.com/technetwork/server-storage/solaris11/downloads/index.html
Oracle Solaris 11 product documentation	http://www.oracle.com/technetwork/server-storage/solaris11/documentation/index.html
Oracle Solaris 11 How-to guides	http://www.oracle.com/technetwork/server-storage/solaris11/documentation/how-to-517481.html
Oracle Solaris 11 training and support	http://www.oracle.com/technetwork/server-storage/solaris11/training/index.html
Official Oracle Solaris blog	https://blogs.oracle.com/solaris/

