

Implement iSCSI solution in Solaris10

Tim Chung

Version 1.2 (OCT, 2009)

QSAN Technology, Inc.
<http://www.qsan.com.tw>
White Paper# **QWP200906-P210C**

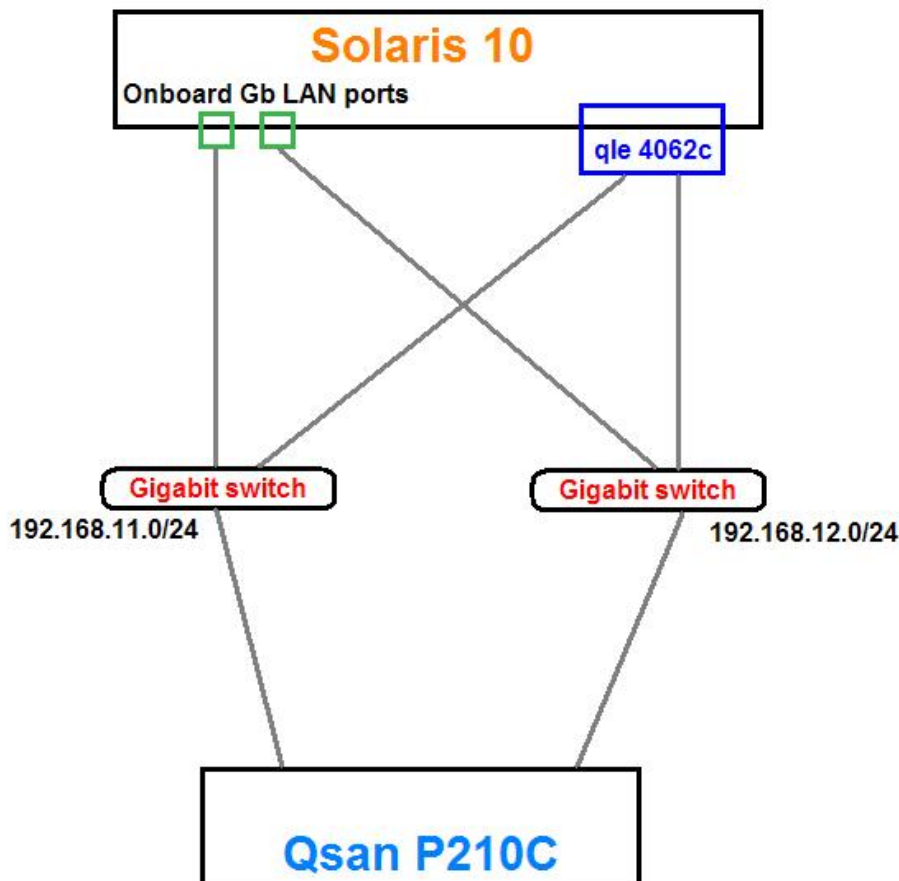
Introduction

In this document, it describes how to connect an iSCSI target in Solaris 10 via software iSCSI initiator and **QLogic QLE4062C** iSCSI HBA. Both of them can work properly with **QSAN P series** models. But only the software iSCSI initiator solution supports MPIO in Solaris 10. In the end, there are performance test results of these two solutions with **QSAN P210C**.

Environment

Server Model: SunFire x2100
iSCSI HBA QLogic QLE4062C
Host OS: Solaris 10 update6 x86 (10/08)
Necessary package: SUNWiscsir, SUNWiscsiu
iSCSI target: QSAN P210C
RAM: 1GB DDR2-667
Firmware: 1.0.3p1 (20081212_1700)
iSCSI data port: 192.168.11.112/24, 192.168.12.112/24

Diagram



Installation

Part 1: QLogic QLE4062C

1. Please look at **QLogic** website first. (There are some useful links in Reference.) Please follow **Quick Start Guide** to install the HBA hardware, download the proper driver and SANsurfer management tool. In order to install the driver, please follow the command:

```
# gzip -dv <directory>/qla4xxx.i386.z
# pkgadd -d <directory>/qla4xxx.i386
# reboot -- -r                                     // Reboot the system to
                                                    // make it take effect.
```

2. To install the SANsurfer management tool, please follow the command:

```
# gzip -dv <directory>/iSCSI_SANsurfer_xxx.tgz
# tar -xvf <directory>/iSCSI_SANsurfer_xxx.tar
# <directory>/iSCSI_SANsurfer_*.bin                // Follow the on-screen
                                                    // instructions to finish
                                                    // the installation.
```

Part 2: Software iSCSI initiator

1. Verify the iSCSI software packages are installed.

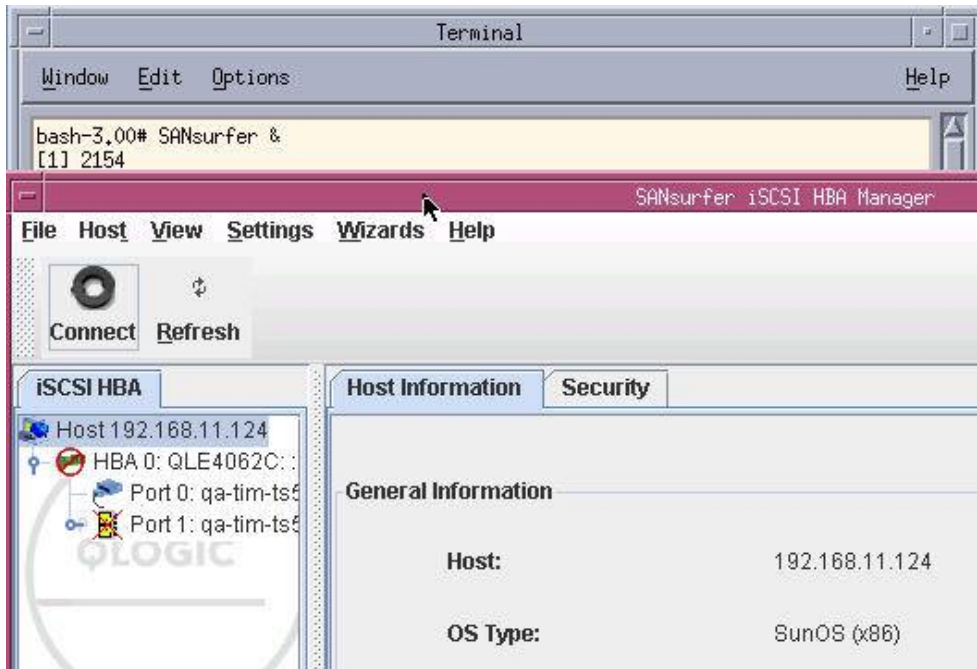
```
# pkginfo SUNWiscsiu SUNWiscsir
system      SUNWiscsiu  Sun   iSCSI Device      Driver      (root)
system      SUNWiscsir  Sun   iSCSI Management  Utilities   (root)
```

If the above packages are not appeared, the Solaris installation CD/DVD is needed.

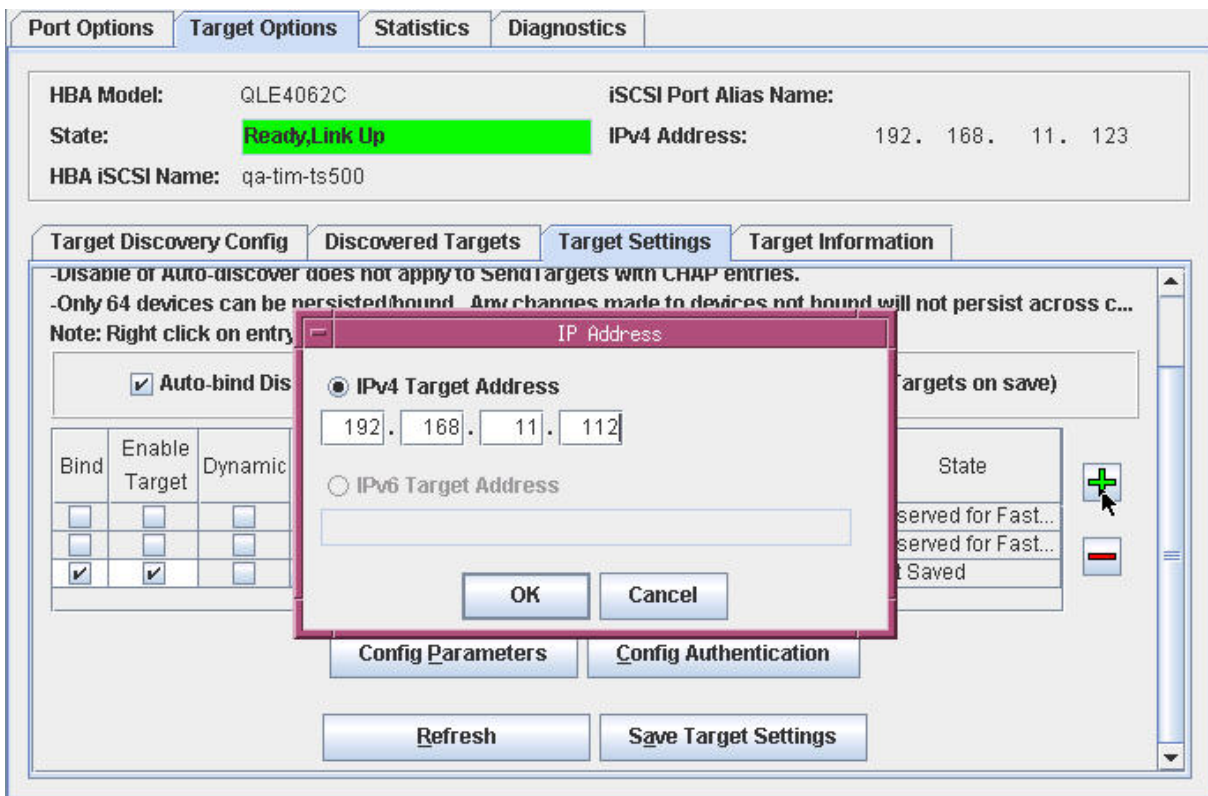
Configuration

Part 1: QLogic QLE4062C

1. Run SANsurfer management tool.



2. Add iSCSI portal as the discovery target of each port on HBA.



3. Save the changes and refresh the HBA.

Port Options | Target Options | Statistics | Diagnostics

HBA Model: QLE4062C iSCSI Port Alias Name:
 State: **Ready, Link Up, Targets Modified** IPv4 Address: 192. 168. 11. 123
 HBA iSCSI Name: qa-tim-ts500

Target Discovery Config | Discovered Targets | Target Settings | Target Information

-Disable or auto-discover does not apply to SendTargets with CHAP entries.
 -Only 64 devices can be persisted/bound. Any changes made to devices not bound will not persist across c...
 Note: Right click on entry for additional features.

Auto-bind Discovered Targets Auto-discover (Re-discover prior SendTargets on save)

Bind	Enable Target	Dynamic	IP Address	iSCSI Name	Target ID	Alias	State
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0.0.0.0	NA	0	NA	Reserved for Fast...
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0.0.0.0	NA	1	NA	Reserved for Fast...
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	192.168.11.112		2		Not Saved

Buttons: Config Parameters, Config Authentication, Refresh, Save Target Settings

4. Refreshing HBA settings will let the iSCSI target connect automatically.

Port Options | Target Options | Statistics | Diagnostics

HBA Model: QLE4062C iSCSI Port Alias Name:
 State: **Ready, Link Up** IPv4 Address: 192. 168. 11. 123
 HBA iSCSI Name: qa-tim-ts500

Target Discovery Config | Discovered Targets | Target Settings | Target Information

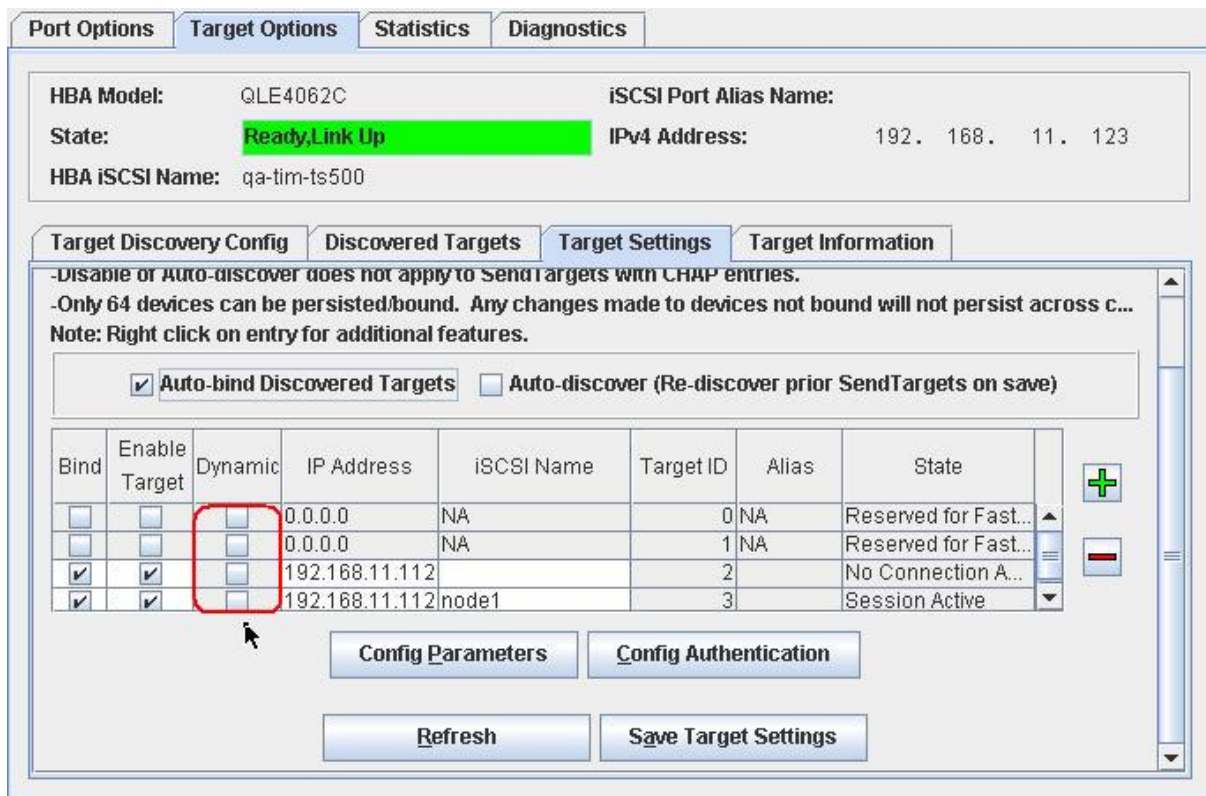
-Disable or auto-discover does not apply to SendTargets with CHAP entries.
 -Only 64 devices can be persisted/bound. Any changes made to devices not bound will not persist across c...
 Note: Right click on entry for additional features.

Auto-bind Discovered Targets Auto-discover (Re-discover prior SendTargets on save)

Bind	Enable Target	Dynamic	IP Address	iSCSI Name	Target ID	Alias	State
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0.0.0.0	NA	0	NA	Reserved for Fast...
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0.0.0.0	NA	1	NA	Reserved for Fast...
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	192.168.11.112		2		No Connection A...
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	192.168.11.112	node1	64		Session Active

Buttons: Config Parameters, Config Authentication, Refresh, Save Target Settings

5. The HBA detects a **dynamic** target. Make sure to save the discover results again to let the targets as **persistent**, otherwise the iSCSI disks will not work properly in Solaris.



HBA Model: QLE4062C iSCSI Port Alias Name:

State: **Ready, Link Up** IPv4 Address: 192. 168. 11. 123

HBA iSCSI Name: qa-tim-ts500

Target Discovery Config Discovered Targets Target Settings Target Information

-Disable or auto-discover does not apply to SendTargets with CHAP entries.
 -Only 64 devices can be persisted/bound. Any changes made to devices not bound will not persist across c...
 Note: Right click on entry for additional features.

Auto-bind Discovered Targets Auto-discover (Re-discover prior SendTargets on save)

Bind	Enable Target	Dynamic	IP Address	iSCSI Name	Target ID	Alias	State
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0.0.0.0	NA	0	NA	Reserved for Fast...
<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	0.0.0.0	NA	1	NA	Reserved for Fast...
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	192.168.11.112		2		No Connection A...
<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	192.168.11.112	node1	3		Session Active

Config Parameters Config Authentication

Refresh Save Target Settings

6. Reboot Solaris to let the newly added LUNs be found.
7. Run the following command after reboot.

```
# grep qla4xxx /var/adm/messages
...
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd10 at qla4xxx target2
lun0
...
```

The message indicates the SCSI device entry which is created in **/kernel/drv/sd.conf** file. This file must be set properly; otherwise the iSCSI target would not be detected. Here are two examples for reference.

Example 1: Connect to **P210C** with two IP portals shared the same nodename. Attach to 3 LUNs which are LUN 0, 1, 254.

After connecting the iSCSI target with QLogic QLE4062C, the messages are displayed like the following.


```
# grep qia4xxx /var/adm/messages
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd10 at qia4xxx target2
lun0
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd11 at qia4xxx target2
lun1
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd12 at qia4xxx target2
lun254
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd13 at qia4xxx target3
lun0
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd14 at qia4xxx target3
lun1
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd15 at qia4xxx target3
lun254
```

Because there are two different IP portals with the same iSCSI node, total 6 LUNs will be found in 2 targets. The configuration file **/kernel/drv/sd.conf** should be added the SCSI information as the following.

```
# vi /kernel/drv/sd.conf
...
name="sd" class="scsi" target=0 lun=0;
name="sd" class="scsi" target=1 lun=0;
name="sd" class="scsi" target=2 lun=0;
name="sd" class="scsi" target=2 lun=1;
name="sd" class="scsi" target=2 lun=254;
name="sd" class="scsi" target=3 lun=0;
name="sd" class="scsi" target=3 lun=1;
name="sd" class="scsi" target=3 lun=254;
name="sd" class="scsi" target=4 lun=0;
name="sd" class="scsi" target=5 lun=0;
...
```

Reboot Solaris to let the settings work, then the disks will be detected.

Example 2: Connect to **P210C** with two IP portals which are shared 3 nodenames, and LUN 0 is attached to node1, LUN 1 is attached to node2, LUN 254 is attached to node3.

After connecting the iSCSI target with QLogic QLE4062C, the messages are displayed like the following.

```
# grep qia4xxx /var/adm/messages
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd10 at qia4xxx target2
lun0
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd11 at qia4xxx target3
lun1
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd12 at qia4xxx target4
lun254
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd13 at qia4xxx target5
lun0
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd14 at qia4xxx target6
lun1
Feb 10 10:30:19 hostname:[ID xxxxxx kernel.info] sd15 at qia4xxx target7
lun254
```

Because there are two different IP portals which are shared 3 iSCSI nodes, total 6 LUNs will be found in 6 targets. The configuration file **/kernel/drv/sd.conf** should be added the SCSI information as the following.

```
# vi /kernel/drv/sd.conf
...
name="sd" class="scsi" target=0 lun=0;
name="sd" class="scsi" target=1 lun=0;
name="sd" class="scsi" target=2 lun=0;
name="sd" class="scsi" target=3 lun=0;
name="sd" class="scsi" target=3 lun=1;
name="sd" class="scsi" target=4 lun=0;
name="sd" class="scsi" target=4 lun=254;
name="sd" class="scsi" target=5 lun=0;
name="sd" class="scsi" target=6 lun=0;
name="sd" class="scsi" target=6 lun=1;
name="sd" class="scsi" target=7 lun=0;
name="sd" class="scsi" target=7 lun=254;
name="sd" class="scsi" target=8 lun=0;
name="sd" class="scsi" target=9 lun=0;
...
```

Reboot Solaris to let the settings work, then the disks will be detected.

8. Reboot Solaris to let the new disk setting validate, and use **devfsadm** command to configure **/dev** namespace of new devices.

```
# devfsadm -Cv // Cleanup unused device
// names.
# devfsadm -i qla4xxx // configure the device
// names for the specified
// driver
```

Part 2: Software iSCSI initiator

1. Configure the device to be discovered dynamically (SendTargets).

```
# iscsiadm add discovery-address 192.168.11.112:3260
```

2. Enable SendTargets discovery method.

```
# iscsiadm modify discovery --sendtargets enable
```

3. Create the iSCSI device links in the **/dev** namespace.

```
# devfsadm -i iscsi
```

4. **(Optional)** Configure CHAP authentication for iSCSI initiator. Use the following commands to set up CHAP secret key. Note that the length of CHAP secret must be during 12 to 16 characters.

```
# iscsiadm modify initiator-node --CHAP-secret
```

5. **(Optional)** Set CHAP name of initiator.

```
# iscsiadm modify initiator-node --CHAP-name your-CHAP-name
```

6. **(Optional)** Enable CHAP authentication.


```
# iscsiadm modify initiator-node -authentication CHAP
```

7. To remove the iSCSI SendTarget discovery entry.

```
# iscsiadm remove discovery-address 192.168.11.112:3260
```

8. To disable the SendTarget discovery method.

```
# iscsiadm modify discovery --sendtargets disable
```

9. Make sure to modify `/kernel/drv/sd.conf` file which is based on LUN settings as the above examples.

Part 3: How to setup multipath driver – MPxIO in Solaris10

1. Verify that the MPxIO configuration parameter is enabled in `/kernel/drv/iscsi.conf` file.

```
# grep mpxio /kernel/drv/iscsi.conf
iscsi.conf: mpxio-disable="no" // "no" means enabled, and
// "yes" means disabled
```

2. Check the configured session parameter.

```
# iscsiadm list initiator-node
Initiator node name: iqn.1986-03.com.sun:01:badad8f3ffff.49af24c2
Initiator node alias: -
  Login Parameters (Default/Configured):
    Header Digest: NONE/-
    Data Digest: NONE/-
  Authentication Type: CHAP
    CHAP Name: qsan
  RADIUS Server: NONE
  RADIUS access: unknown
  Configured Sessions: 1
```

3. Set the configured session parameter to 2. The value must be during 1 to 4.

```
# iscsiadm modify initiator-node -c 2
```

Or configured session can be bound more than one local IP address, like this:

```
# iscsiadm modify initiator-node -c 192.168.11.1, 192.168.12.1
```

4. Verify that the parameter has been modified.

```
# iscsiadm list initiator-node
...
  Configured Sessions: 2
```

- Next, tell MPxIO driver to accept the device based on the SCSI vendor and Product-ID by editing **/kernel/drv/scsi_vhci.conf** file. Add the following settings in the tail of the file.

```
# vi /kernel/drv/scsi_vhci.conf
...
device-type-scsi-options-list=
"Qsan    P210C", "symmetric-option";           // 4 spaces must be kept
                                                // between Qsan and
symmetric-option = x01000000;                 // P210C
```

In order to find out the SCSI vendor and Product-ID, use **format** command and select one of the instances in device.

```
format> inquiry
Vendor:      Qsan
Product:     P210C
Revision:    1.0.
```

Beware that the MPxIO driver accepts the vendor name as 8 character-words only. So the vendor name must be added 4 spaces at the end of "Qsan" in **/kernel/drv/scsi_vhci.conf** file.

- A reboot is necessary to let all settings validate.

```
# reboot -- -r
```

- After rebooting, these iSCSI disks will be merged to one MPIO device.

```
# format // Before MPxIO applied,
// OS detects the
sameSearching for disks... // LUN as two independent
// disks.
...
AVAILABLE DISK SELECTIONS:
  0. c1t1d0 <ATA-HDS728040PLA320-A69A cyl 5003 alt 2 hd 255 sec 63>
    /pci@0,0/pci108e,5348@7/disk@1,0
  2. c3t266d0 <DEFAULT cyl 1303 alt 2 hd 255 sec 63>
    /iscsi/disk@0000node10000,0
  3. c3t267d0 <DEFAULT cyl 1303 alt 2 hd 255 sec 63>
    /iscsi/disk@0000node10000,0
Specify disk (enter its number):
```

```
# format // After MPxIO applied,
// two disks c3t266d0 and
Searching for disks... // c3t267d0 are merged to
// c4t2436001378A6D1B6d0
...
AVAILABLE DISK SELECTIONS:
  0. c1t1d0 <ATA-HDS728040PLA320-A69A cyl 5003 alt 2 hd 255 sec 63>
    /pci@0,0/pci108e,5348@7/disk@1,0
  1. c4t2436001378A6D1B6d0 <DEFAULT cyl 1303 alt 2 hd 255 sec 63>
    /scsi_vhci/disk@g2436001378a6d1b6
Specify disk (enter its number):
```

Performance test

	SW iSCSI initiator		QLogic QLE4062C	
Read	I/Ops: 60358	CPU util: 90%	I/Ops: 36585	CPU util: 65%
	MBps: 111	CPU util: 33%	MBps: 112	CPU util: 30%
Write	I/Ops: 73647	CPU util: 97%	I/Ops: 34100	CPU util: 62%
	MBps: 112	CPU util: 29%	MBps: 112	CPU util: 27%

Test tool:

lometer 2006.07.27

Test configuration:

512B, 100% Read, 0% Random; 256K, 100% Read, 0% Random.

512B, 0% Read, 0% Random; 256K, 0% Read, 0% Random.

RAID setting:

RAID0 with 8 SATA disks, 50GB UDV (Stripe 64K, Block size 512B)

Summary

In this document, two types of iSCSI initiator are demonstrated. With QLogic QLE4062C iSCSI HBA, it is easily configured by SANsurfer management tool. But it can not cooperate with MPxIO driver in Solaris to provide the feature of MPIO. While the software iSCSI initiator is inconvenient to setup, it can work with the MPxIO driver to provide a reliable MPIO mechanism.

Applies to

- All **QSAN P series** controllers FW (20081212_1700)

References

- QLogic QLE4062C driver for Solaris x86
http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/SearchByProduct.aspx?ProductCategory=82&Product=1037&Os=145
- Quick Start Guide, iSCSI HBA Installation
http://filedownloads.qlogic.com/files/Manual/60311/QuickStartGuide_iSCSIHBAs_C.pdf
- SANsurfer iSCSI Manager User's Guide
<http://filedownloads.qlogic.com/files/Manual/69200/IS0054602-00B.pdf>
- Configuring Solaris iSCSI Target and Initiators
<http://docs.sun.com/app/docs/doc/817-5093/fmvcd?a=view>
- Solaris 10 with MPxIO (scsi_vhci.conf)...
<http://southbrain.com/south/2008/02/solaris-10-with-mpxio-scsi-vhc.html>