



Solaris 10 Best Practice with NeoSapphire P710

Table of Content

1. Introduction.....	3
2. Verifying FC adapter information.....	4
3. Configuring MPxIO.....	5
4. Formatting a volume configuration on Solaris 10	7
5. Performance Tuning.....	10
6. References	12

1. Introduction

This document provides a comprehensive guide to configuring the NeoSapphire P710 Fibre Channel (P710-FC) volumes within Solaris 10.

The following steps are covered in this guide:

1. Verifying the FC HBA information.
2. Identifying the partition of a FC disk.
3. Configuring multipath for Solaris 10.
4. Performance tuning tips for P710-FC within Solaris 10.

Hardware Configuration:

The lab scenario in this example was performed using a freshly installed instance of Solaris 10, update 1/13 with the following hardware configuration:

1. CPU: Intel Xeon® E5-2690 v4
2. Memory: 64GB DDR4-2400
3. Local SSD: 2.5" 800GB SATA3 SSD drive
4. Two port 16G Fibre Channel HBA: Emulex LPE16002B-M6

FC Topology:

The server has both FC ports connected to a Brocade 6506 16Gbps FC switch, and from the FC switch, the P710-FC is connected via 6 FC cables to the 6 FC ports on the P710-FC.

2. Verifying FC adapter information

After completing the Solaris 10 installation, the emlxs driver for the Emulex 16G FC HBA was automatically installed and configured. When using Emulex FC HBA for FC connections, make sure you to use the emlxs driver instead of the original lpfc driver.

In Solaris 10, type `fcinfo hba-port` to verify the existing FC HBA port status:

```
bash-3.2# fcinfo hba-port
HBA Port WWN: 10000090faf05c26
  OS Device Name: /dev/cfg/c2
  Manufacturer: Emulex
  Model: LPe16002B-M6
  Firmware Version: 11.1.217.0 (xe201.grp)
  FCode/BIOS Version: Boot:11.1.217.0
  Serial Number: FC70668585
  Driver Name: emlxs
  Driver Version: 2.80.8.0 (2012.09.17.15.45)
  Type: N-port
  State: online
  Supported Speeds: 4Gb 8Gb 16Gb
  Current Speed: 16Gb
  Node WWN: 20000090faf05c26
HBA Port WWN: 10000090faf05c27
  OS Device Name: /dev/cfg/c3
  Manufacturer: Emulex
  Model: LPe16002B-M6
  Firmware Version: 11.1.217.0 (xe201.grp)
  FCode/BIOS Version: Boot:11.1.217.0
  Serial Number: FC70668585
  Driver Name: emlxs
  Driver Version: 2.80.8.0 (2012.09.17.15.45)
  Type: N-port
  State: online
  Supported Speeds: 4Gb 8Gb 16Gb
  Current Speed: 16Gb
  Node WWN: 20000090faf05c27
```

Here we can see there were two FC ports successfully identified and connected at 16Gbps speed.

3. Configuring MPxIO

Before formatting the disk device in Solaris 10, the disk must be properly configured with MPxIO. In this example we used 2 paths from the client to a Fibre Channel switch, and then 6 paths from the switch to the P710. There were initially 12 devices detected before MPxIO was properly configured for Solaris.

Type format and all the currently detected/connected fibre channel disks will be listed as below (without mpzio configuration): (Use Ctl+C to interrupt and quit here).

```
AVAILABLE DISK SELECTIONS:
 0. clt0d0 <ATA      -2.5" SATA SSD 3-114  cyl 45678 alt 2 hd 224 sec 168>
    /pci@0,0/pci15d9,87d@1f,2/disk@0,0
 1. c2t10000090FAF05B16d0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0/fp@0,0/disk@w10000090faf05b16,0
 2. c2t10000090FAF05B17d0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0/fp@0,0/disk@w10000090faf05b17,0
 3. c2t100000109B17B86Ed0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0/fp@0,0/disk@w100000109b17b86e,0
 4. c2t100000109B17B86Fd0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0/fp@0,0/disk@w100000109b17b86f,0
 5. c2t100000109B1799BCd0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0/fp@0,0/disk@w100000109b1799bc,0
 6. c2t100000109B1799BDd0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0/fp@0,0/disk@w100000109b1799bd,0
 7. c3t10000090FAF05B16d0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0,1/fp@0,0/disk@w10000090faf05b16,0
 8. c3t10000090FAF05B17d0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0,1/fp@0,0/disk@w10000090faf05b17,0
 9. c3t100000109B17B86Ed0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0,1/fp@0,0/disk@w100000109b17b86e,0
10. c3t100000109B17B86Fd0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0,1/fp@0,0/disk@w100000109b17b86f,0
11. c3t100000109B1799BCd0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0,1/fp@0,0/disk@w100000109b1799bc,0
12. c3t100000109B1799BDd0 <Astor-P710- 300-4.20TB>
    /pci@0,0/pci8086,6f08@3/pci10df,e280@0,1/fp@0,0/disk@w100000109b1799bd,0
Specify disk (enter its number): ^C
#
```

To enable the mpzio for Solaris, prepare and modify the following files:

1. Modify **/kernel/drv/fp.conf** and make sure the following settings is present in the configuration file, add it if not:

```
mpzio-disable="no";
```

2. Modify **/kernel/drv/scsi_vhci.conf** with the following settings: (For **device-type-scsi-option-list** option, use "Astor" as the vendor ID and "P710" as the product ID in this configuration file.)

```
load-balance="round-robin";
```

```
auto-failback="enable";
```

```
device-type-scsi-options-list ="AStor P710", "symmetric-option";  
symmetric-option = 0x1000000;
```

3. Enable MPxIO service for Fibre Channel only with the following command:

```
# stmsboot -D fp -e
```

(A Reboot here is necessary for the mpio configurations to take effect, the system will automatically reboot twice).

* MPxIO can be globally enabled with **#stmsboot -e**

* MPxIO can be globally disabled with **#stmsboot -d**

4. Verify the status of MPxIO by checking scsi_vhci module status with the following command:

```
# modinfo -c |grep scsi_vhci
```

5. Successful configuration of MPxIO will merge all FC devices under a scsi_vhci path as follows:

```
1. c4t34074FC1F2A8291Cd0 <AStor -P710 - 300 cyl 23403 alt 2 hd 224 sec 56>  
 /scsi_vhci/disk@g34074fc1f2a8291c  
2. c4t34074FC17A560FBA0 <AStor -P710 - 300 cyl 3110 alt 2 hd 224 sec 616> P710test  
 /scsi_vhci/disk@g34074fc17a560fba
```

(This screenshot shows there were two volumes present, both merged to **/scsi_vhci/disk@xxxxxxxx** paths.)

6. Verify the health status of multipath with the following command:

```
# mpathadm list LU
```

```
/dev/rdisk/c4t34074FC1756EC02Ed0s2
```

```
Total Path Count: 12
```

```
Operational Path Count: 12
```

The number of paths detected and displayed here varies based on your FC topology.

7. The target disk is then ready to be formatted.

4. Formatting a volume configuration on Solaris 10

Use format command to list current detected disk device, and select the target disk to be formatted.

```
# format -e
Searching for disks...done

AVAILABLE DISK SELECTIONS:
    0. c1t0d0 <ATA      -2.5" SATA SSD 3-114   cyl 45678 alt 2 hd 224 sec
168>
        /pci@0,0/pci15d9,87d@1f,2/disk@0,0
    1. c4t34074FC1F2A8291Cd0 <AStor  -P710           - 300 cyl 23402
alt 2 hd 224 sec 56>
        /scsi_vhci/disk@g34074fc1f2a8291c
    2. c4t34074FC17A560FBA0 <AStor-P710- 300-4.20TB>
        /scsi_vhci/disk@g34074fc17a560fba

Specify disk (enter its number): 1
selecting c4t34074FC1F2A8291Cd0
[disk formatted]
```

Then we'll need to re-partition the select disk and adjust the 4K alignment for the disk. By default a Solaris 10 formatted disk does not partition with 4K aligned offset. This causes a performance penalty and we need to fix it using the steps below.

```
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[0]: 1
Warning: This disk has an SMI label. Changing to EFI label will erase all
current partitions.
Continue? yes
format> partition
.....
partition> print
Current partition table (default):
Total disk sectors available: 293584861 + 16384 (reserved sectors)
Part      Tag      Flag      First Sector      Size      Last Sector
```

```

0          usr      wm                34          139.99GB
293584861
 1 unassigned  wm                0           0           0
 2 unassigned  wm                0           0           0
 3 unassigned  wm                0           0           0
 4 unassigned  wm                0           0           0
 5 unassigned  wm                0           0           0
 6 unassigned  wm                0           0           0
 7 unassigned  wm                0           0           0
 8 reserved    wm                293584862    8.00MB
293601245
partition> 0
Part      Tag      Flag      First Sector      Size      Last Sector
 0          usr      wm                34          139.99GB
293584861
Enter partition id tag[usr]:
Enter partition permission flags[wm]:
Enter new starting Sector[34]: 256
Enter partition size[293584828b, 293585083e, 143351mb, 139gb, 0tb]:
293584861e
partition> print
Current partition table (unnamed):
Total disk sectors available: 293584861 + 16384 (reserved sectors)
Part      Tag      Flag      First Sector      Size      Last Sector
 0          usr      wm                256          139.99GB
293584861
 1 unassigned  wm                0           0           0
 2 unassigned  wm                0           0           0
 3 unassigned  wm                0           0           0
 4 unassigned  wm                0           0           0
 5 unassigned  wm                0           0           0
 6 unassigned  wm                0           0           0
 7 unassigned  wm                0           0           0
 8 reserved    wm                293584862    8.00MB
293601245
partition> label
[0] SMI Label
[1] EFI Label

```

```
Specify Label type[1]: 1
Ready to label disk, continue? yes
partition> quit
.....
format> quit
#
```

In the steps above we re-partitioned the selected disk (disk 1), to use a starting sector size of 256KB instead of the original 34KB, which is not 4K aligned. For a disk size of more than 2TB the disk needs to be labelled with an EFI label. During the configuration the disk size needs to be entered with the result printed from “print” command and modified. For example, in this case the printed result was **293584862**, we used this value and subtract 1 then append an “e” as the input value **293584861e** for the “Enter partition size” prompt. At the end the disk’s partition table was successfully modified to a 4KB aligned configuration with an unchanged partition size. The disk will need to be re-labelled again since the initial label information will be lost by manually changing the partition information.

After the disk is formatted and partitioned, select a partition to make new file system with. In Solaris 10 the partition 0~8 as listed in the partition results in previous steps can be accessed via `/dev/dsk/c4t34074FC17A560FBAd0s0` to**s8** as in Solaris 10 the partitions are presented as “slices” in every disk device. Use **newfs** to format the target partition and then mount the partition with UFS file system format (with parameter **-F ufs**).

```
newfs /dev/dsk/c4t34074FC17A560FBAd0s0
mount -F ufs /dev/dsk/c4t34074FC17A560FBAd0s0 /mnt/test
```

5. Performance Tuning

In this example we used the following modifications to adjust the P710-FC performance.

Modify **/etc/system** with

```
set sd:sd_io_time = 0x3C
set maxphys = 0x400000
```

Modify **/kernel/drv/sd.conf** with

```
sd_max_xfer_size=0x400000;
sd-config-list="AStor P710", "physical-block-size:4096, retries-busy:30, retries-
reset:30, retries-notready:300, retries-timeout:10, throttle-max:64, throttle-
min:8";
```

Modify **/kernel/drv/emlxs.conf** with

```
target-depth=512
```

The **symmetric-option** must be set properly and double verified to take effect. Make sure with the following command the Asymmetric option is displayed as “no”:

```
bash-3.2# mpathadm show lu /dev/rdisk/c4t34074FC1756EC02Ed0s2
Logical Unit: /dev/rdisk/c4t34074FC1756EC02Ed0s2
    mpath-support: libmpscsi_vhci.so
    Vendor: AStor
    Product: P710
    Revision: 300
    Name Type: unknown type
    Name: 34074fc1756ec02e
    Asymmetric: no
    Current Load Balance: round-robin
    Logical Unit Group ID: NA
    Auto Failback: on
    Auto Probing: NA

    Paths:
        Initiator Port Name: 10000090faf05c26
        Target Port Name: 100000109b1799bd
```

Override Path: NA

Path State: OK

Disabled: no

.....

Target Ports:

Name: 100000109b17b86e

Relative ID: 0

.....

6. References

EFI partitions / GUID partition table:

<http://www.oracle.com/technetwork/server-storage/sun-unified-storage/documentation/partitionalign-111512-1875560.pdf>

Solaris 10 version check:

<https://community.oracle.com/thread/2244132?start=0>

MPxIO configuration:

<http://docs.oracle.com/cd/E19253-01/820-1931/6ndinfi7p/index.html>