

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1



Purpose: This document describes how to implement SUN's MPXIO path manager to allocate AMS/WMS LUNs across both controllers of AMS/WMS using Solaris command mpathadm. Once MPXIO and AMS/WMS is configured as described below then both Controllers of an AMS/WMS will be utilized. As in all AMS/WMS multipath implementations MPXIO can not be used for Load Balancing of I/O to LUNs across the 2 AMS/WMS controllers.

Prerequisites:

- Solaris 10 U3 SPARC or above installed. Solaris 10 X64 currently not supported using this technique.
- Solaris bundled HBA driver.

Warnings:

-If an intermittent path failure exists and the customer executes the script while the path is not stable, it is expected that applications running on the server will receive an error.

-Execution of the script is necessary after rebooting the OS and recovering from path failures.

Setup:

Step 1. Setup connectivity, zones etc.

Step 2. Configure Hitachi AMS/WMS storage Host Group Options and as necessary , the Host connection mode 2 Options as in the examples below using AMS/WMS Storage Navigator:

AMS/WMS series setting

Host Group Options

Platform:	Solaris
Alternate Path:	MPxIO
Failover:	None or Sun Cluster
Host Connection Mode 1:	Standard Mode
Host Connection Mode 2:	None (*1)

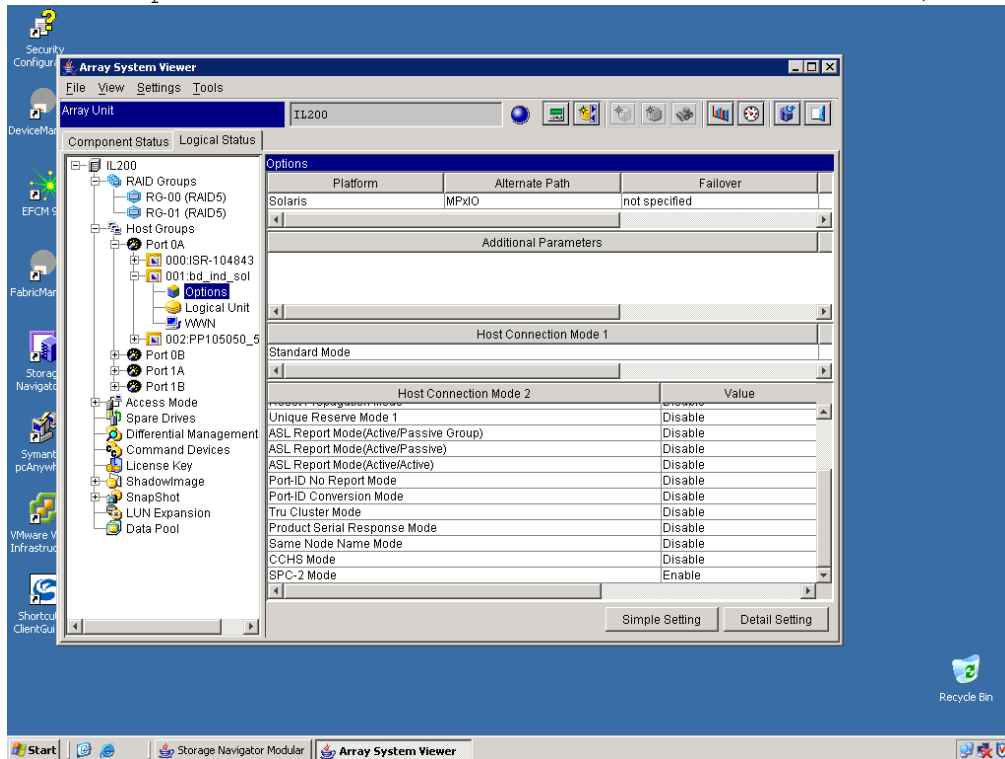
(*1) When you use Solaris Volume Manager(SVM), please enable the following options.

- SPC-2 Mode
- Reset Propagation Mode

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

The Following is an example of the Storage Navigator screen which illustrates the options currently enabled (Reset Propagation and SPC-2 are options are not shown in this case - No SVM case)



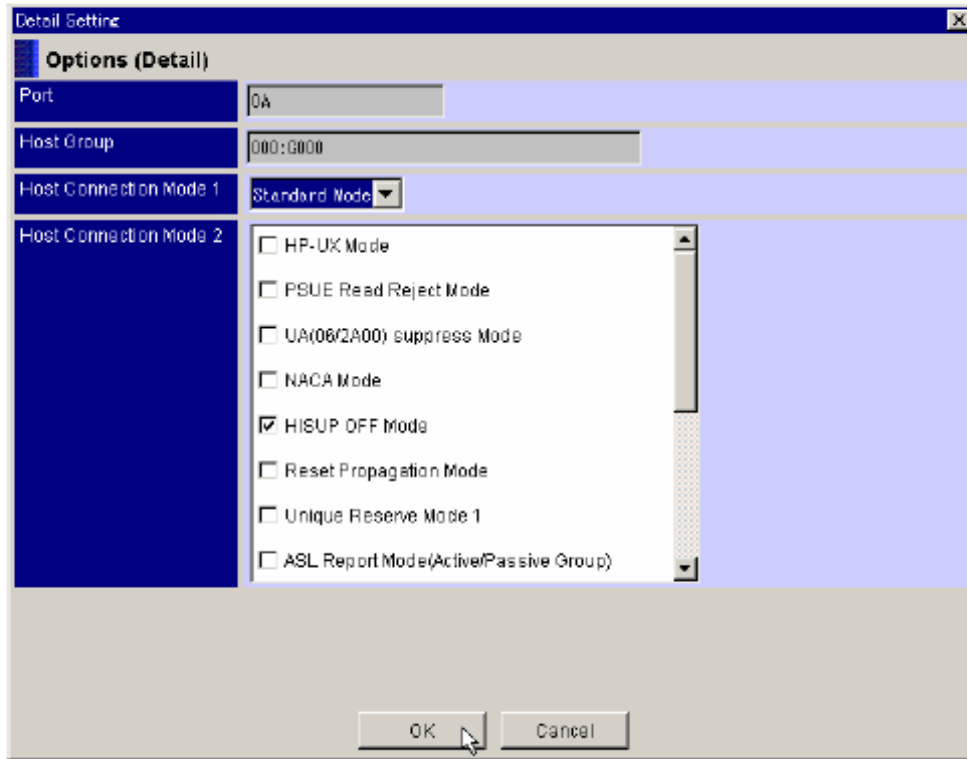
Example 1: AMS options settings

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

The following is an example of the Storage Navigator section which is used to change the Host Connection mode2 options for the SVM case (Reset propagation and SPC-2)

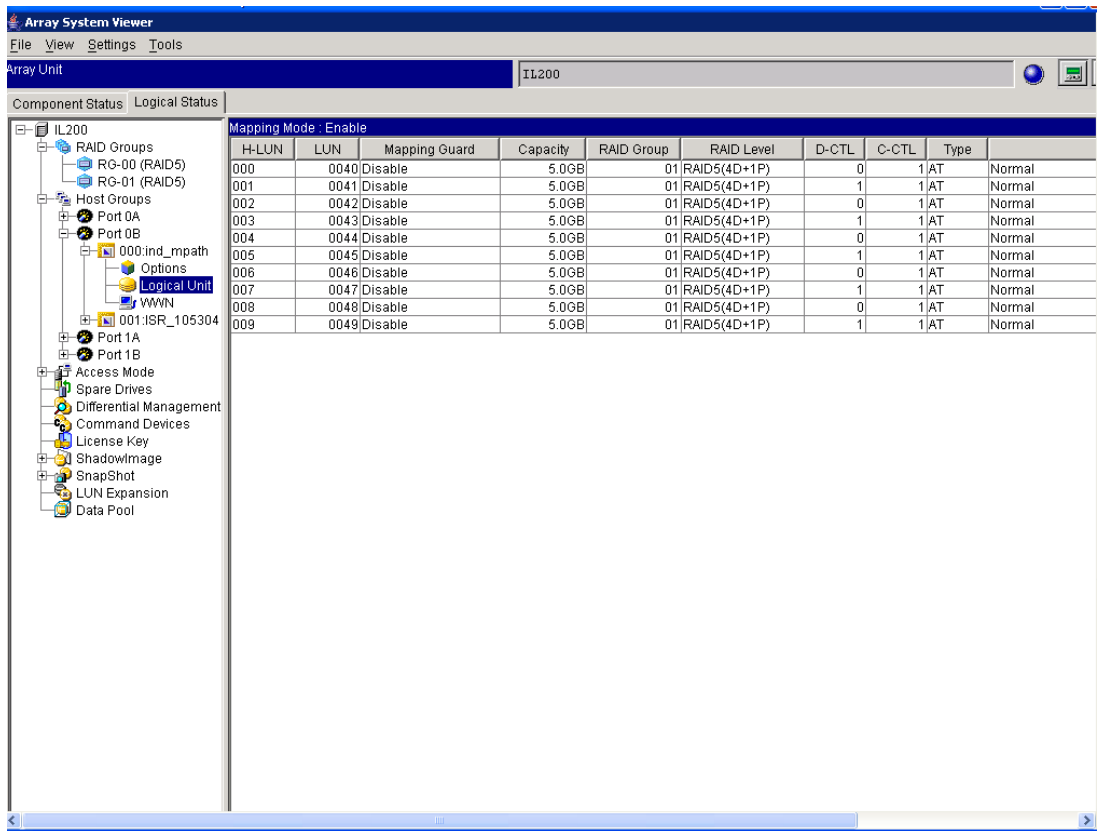
9. Select the Detail button.



IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

Step 3. The LUNs are assigned in such a way that default controller(D-CTL) of some LUNs is CTL-0 and default controller(D-CTL) for the other LUNs is CTL-1.



The screenshot shows the Array System Viewer interface for IL200. The 'Logical Status' tab is active, displaying a tree view on the left and a table of LUNs on the right. The table lists LUNs 000 through 009, their capacities (all 5.0GB), RAID groups (01), RAID levels (RAID5(4D+1P)), and D-CTL/C-CTL assignments. LUNs 000, 002, 004, 006, and 008 have D-CTL 0 and C-CTL 1. LUNs 001, 003, 005, 007, and 009 have D-CTL 1 and C-CTL 0.

H-LUN	LUN	Mapping Guard	Capacity	RAID Group	RAID Level	D-CTL	C-CTL	Type
000	0040	Disable	5.0GB	01	RAID5(4D+1P)	0	1	AT Normal
001	0041	Disable	5.0GB	01	RAID5(4D+1P)	1	1	AT Normal
002	0042	Disable	5.0GB	01	RAID5(4D+1P)	0	1	AT Normal
003	0043	Disable	5.0GB	01	RAID5(4D+1P)	1	1	AT Normal
004	0044	Disable	5.0GB	01	RAID5(4D+1P)	0	1	AT Normal
005	0045	Disable	5.0GB	01	RAID5(4D+1P)	1	1	AT Normal
006	0046	Disable	5.0GB	01	RAID5(4D+1P)	0	1	AT Normal
007	0047	Disable	5.0GB	01	RAID5(4D+1P)	1	1	AT Normal
008	0048	Disable	5.0GB	01	RAID5(4D+1P)	0	1	AT Normal
009	0049	Disable	5.0GB	01	RAID5(4D+1P)	1	1	AT Normal

Example 2: LUN assignment.

In the above example D-CTL or owner controller of LUN 0,2,4,6,8 is controller 0 and for other LUNs it is controller 1. D-CTL is specified during LUN creation from RAID Group. However, it can be changed later. C-CTL of a LUN denotes the controller through which the LUN is accessed or I/O is going through currently.

Step 4. Enable MPxIO

Multipathing is disabled by default for FC devices on SPARC based systems.

To enable multipathing (MPxIO) perform the following steps:

a) Edit /kernel/drv/scsi_vhci.conf file and set as below

```
(i) load-balance="none";
(ii) auto-failback="enable";
(iii) device-type-scsi-options-list =
      "HITACHI DF600F", "symmetric-option";

symmetric-option = 0x1000000;
```

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

b) Run the following command: **# stmsboot -e**

```
WARNING: This operation will require a reboot.
Do you want to continue ? [y/n] (default: y) y
The changes will come into effect after rebooting the system.
Reboot the system now ? [y/n] (default: y) y
```

Note: To disable MPxIO, run "**stmsboot -d**" command. This is not required to run in this case.

Step 5. After reboot, run "**cfgadm -c configure <HBA>**" to configure the LUNs.

Step 6. Run following commands to verify devices are configured.

a)

```
#>format
```

```
Searching for disks...done
```

```
AVAILABLE DISK SELECTIONS:
```

0. c15t60060E801029C800047905400000002Ad0 <HITACHI-DF600F-0000 cyl 271 alt 2 hd 50 sec 768>
/scsi_vhci/ssd@g60060e801029c800047905400000002a
1. c15t60060E801029C800047905400000002Bd0 <HITACHI-DF600F-0000 cyl 271 alt 2 hd 50 sec 768>
/scsi_vhci/ssd@g60060e801029c800047905400000002b
2. c15t60060E801029C800047905400000002Cd0 <HITACHI-DF600F-0000 cyl 271 alt 2 hd 50 sec 768>
/scsi_vhci/ssd@g60060e801029c800047905400000002c
3. c15t60060E801029C800047905400000002Dd0 <HITACHI-DF600F-0000 cyl 271 alt 2 hd 50 sec 768>
/scsi_vhci/ssd@g60060e801029c800047905400000002d
4. c15t60060E801029C800047905400000002Ed0 <HITACHI-DF600F-0000 cyl 271 alt 2 hd 50 sec 768>

b)

```
#>luxadm display /dev/rdisk/c15t60060E801029C800047905400000002Ad0s2
```

```
DEVICE PROPERTIES for disk: /dev/rdisk/c15t60060E801029C800047905400000002Ad0s2
```

```
Vendor:          HITACHI
Product ID:      DF600F
Revision:        0000
Serial Num:      75040064002A
Unformatted capacity: 5120.000 MBytes
Read Cache:      Enabled
  Minimum prefetch: 0x0
  Maximum prefetch: 0x0
Device Type:     Disk device
Path(s):
```

```
/dev/rdisk/c15t60060E801029C800047905400000002Ad0s2
/devices/scsi_vhci/ssd@g60060e801029c800047905400000002a:c,raw
Controller       /devices/pci@9,700000/SUNW,emlxs@4/fp@0,0
Device Address   50060e801029c801,2
Host controller port WWN 10000000c957ff36
Class            primary
State            ONLINE
Controller       /devices/pci@8,700000/SUNW,emlxs@2/fp@0,0
Device Address   50060e801029c803,2
Host controller port WWN 10000000c95d0005
Class            primary
State            ONLINE
```

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

c)

#>mpathadm list lu

```
/dev/rdisk/c15t20000000879FD5EAd0s2
  Total Path Count: 1
  Operational Path Count: 1
/dev/rdisk/c15t20000000879FDA70d0s2
  Total Path Count: 1
  Operational Path Count: 1
/dev/rdisk/c15t20000000879FD0A4d0s2
  Total Path Count: 1
  Operational Path Count: 1
/dev/rdisk/c15t20000000879F87AAd0s2
  Total Path Count: 1
  Operational Path Count: 1
/dev/rdisk/c15t60060E801029C8000479054000000031d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C8000479054000000030d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C800047905400000002Fd0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C800047905400000002Ed0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C800047905400000002Dd0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C800047905400000002Cd0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C800047905400000002Bd0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C800047905400000002Ad0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C8000479054000000029d0s2
  Total Path Count: 2
  Operational Path Count: 2
/dev/rdisk/c15t60060E801029C8000479054000000028d0s2
  Total Path Count: 2
  Operational Path Count: 2
```

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

How to distribute I/O through owner controller using `mpathadm` command:

The multipath configuration used for the following instructions is a Solaris host with 2 HBA paths, one connected to one FC port of CTL-0 and the other connected to one FC port of CTL-1 of an AMS/WMS storage.

Step 1. Start I/O. Before proceeding to the next steps, I/O should be run to all LUNs using `dd` command to read the raw devices or `cp` command to read/write any data to the file systems.

All of the above LUN's I/O will go through only one of the 2 AMS/WMS controllers. This can be observed by using the Storage Navigator to view the logical LUN status (Example 2 Figure above). This screen will show under column "C-CTL" that all the LUNs current path are controlled on the same controller.

The next steps will show how to have MPXIO distribute LUN I/O across both AMS/WMS controllers by setting preferred controller paths for each of the LUNs.

Step 2. Use `mpathadm` command to disable the path of a LUN whose I/O is going through non-owner controller. This will force its I/O to go through owner controller. Do it for all LUNs having I/O through non-owner controller.

Command format:

```
#mpathadm disable path -i <initiator port name> -t <target name> -l  
<logical unit name>
```

Example:

```
For LUN 0042 (See Example 2 above)  
Device name: /dev/rdisk/c15t60060E801029C800047905400000002Ad0s2 (42=Hex 2A)  
Owner controller WWN           : 50060e801029c801 (CTL-0, Port 0B)  
Non-owner controller WWN       : 50060e801029c803 (CTL-1, Port 1B)  
Initiator WWN to owner controller : 10000000c957ff36  
Initiator WWN to Non-owner controller : 10000000c95d0005  
(See luxadm output above for WWNs)
```

Command to run:

```
mpathadm disable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C800047905400000002Ad0s2
```

Step 3. Enable both paths of all LUNs which were disabled in the previous step. This is to make the alternate path available in case of (owner) path failure.

Example:

```
mpathadm enable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C800047905400000002Ad0s2
```

Now, I/O to all LUNs is routed through respective owner controller and non-owner path is available for failover.

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

Step 4. After each instance of path failure and recovery, step b) and c) must be run to distribute I/O load back to owner controller. This is because, autofailback does not fail back I/O automatically even if auto-failback="enable"; parameter is set in /kernel/drv/scsi_vhci.conf file.

Sample scripts:

1. Disable script-

```
#>cat mpath-dis
```

```
mpathadm disable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C800047905400000028d0s2  
mpathadm disable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Ad0s2  
mpathadm disable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Cd0s2  
mpathadm disable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Ed0s2  
mpathadm disable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C800047905400000030d0s2  
mpathadm disable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C800047905400000029d0s2  
mpathadm disable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Bd0s2  
mpathadm disable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Dd0s2  
mpathadm disable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Fd0s2  
mpathadm disable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C800047905400000031d0s2
```

2. Enable script-

```
#>cat mpath-ena
```

```
mpathadm enable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C800047905400000028d0s2  
mpathadm enable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Ad0s2  
mpathadm enable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Cd0s2  
mpathadm enable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Ed0s2  
mpathadm enable path -i 10000000c95d0005 -t 50060e801029c803 -l  
/dev/rdisk/c15t60060E801029C800047905400000030d0s2  
mpathadm enable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C800047905400000029d0s2  
mpathadm enable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Bd0s2  
mpathadm enable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Dd0s2  
mpathadm enable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C80004790540000002Fd0s2  
mpathadm enable path -i 10000000c957ff36 -t 50060e801029c801 -l  
/dev/rdisk/c15t60060E801029C800047905400000031d0s2
```

IMPLEMENTING MPXIO FOR AMS/WMS

Rev 1

How to automate the script:

The above steps can be automated through /etc/rc script. It is observed that unless I/O is run to a LUN, mpathadm enable or disable command does not show any changes in behavior. A script with the format below can be used as a /etc/rc script S999rc-mpath, under /etc/rc3.d:

- dd if=<device name> of=/dev/null bs=1k
- mpathadm disable <non-owner path>
- mpathadm enable <non-owner path>
- kill dd job

Note:

1. In case of a path failure, all file systems get mounted with available path during system boot. When the path is recovered, the script can be run manually to distribute the I/O.
2. In case of path failure and recovery while server is running, the script must be run manually to distribute I/O load across controllers.