

AIX MPIO for XP Quick Connect Guide

Featuring MPIO version 5.4.0.2

Table of Contents

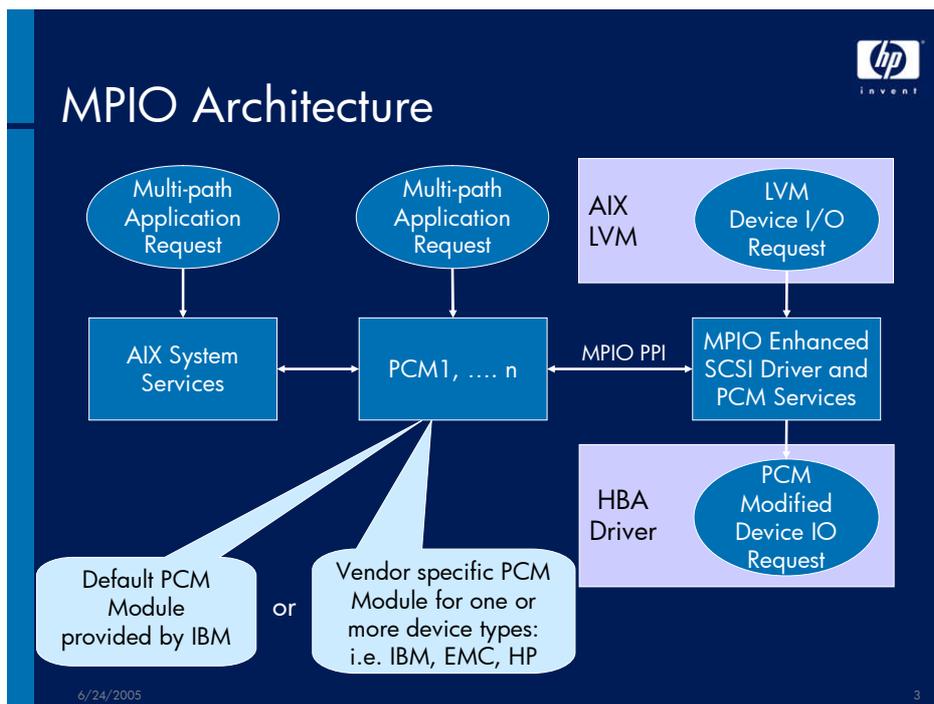
INTRODUCTION	2
SUPPORT	4
SYSTEM REQUIREMENTS	4
SAN BOOT	5
DYNAMIC TRACKING	5
COEXISTENCE	5
<i>Single Path ODM</i>	5
# <i>lslpp -l HP.aix.support.rte</i>	5
<i>HDLM XP (Hitachi Dynamic Link Manager for XP arrays)</i>	6
# <i>lslpp -l DLManager.mpio.rte</i>	6
<i>Other MPIO implementations</i>	6
VIO SERVER	6
MICRO PARTITIONING	6
HACMP	6
RAID MANAGER	7
MPIO FUNCTIONALITY	8
INSTALLATION	10
CHECK FOR XP24000 DEVICES ALREADY CONFIGURED ON THE SYSTEM	10
CHECK FOR THE NON-MPIO XP ODM FILESET	10
INSTALL VIA COMMAND LINE	10
INSTALL VIA SMIT	11
CHECK INSTALLATION	13
DEVICE CONFIGURATION	13
SOFTWARE REMOVAL	16
MPIO OPERATION	18
EXAMPLES	18
<i>lspath</i>	18
<i>chpath</i>	19
<i>rmpath</i>	19
<i>cfgmgr</i>	19
<i>lspv</i>	20
<i>lsdev</i>	20
<i>lscfg</i>	20
<i>lsattr</i>	21
<i>chdev</i>	22
<i>rmdev</i>	22
<i>iostat</i>	23
RELEASE HISTORY	24

Introduction

Starting with AIX 5.2 ML01 and AIX 5.3 ML00, IBM introduced Multiple Path I/O (MPIO) in the AIX base operating system. MPIO is an enhancement to AIX that will provide support for multiple connections from a host to external storage subsystems. A device can be uniquely detected via one or more physical connections, or paths. A path-control module (PCM) provides the path management functions.

AIX MPIO is an enhancement to the base OS environment that will provide support for connection to multi-path Fibre Channel storage subsystems. AIX servers configured with multiple Fibre Channel (FC) Host Bus Adapters (HBA) connected to a Storage Area Network (SAN) can increase the storage subsystem's availability.

To make use of the MPIO architecture for a specific device it needs to be specified as MPIO-capable device in the database of the Object Data Manager (ODM) of AIX. In the ODM the MPIO specific drivers, methods and predefined attributes are specified to support detection, configuration and management of multiple paths for this device.



The MPIO architecture provides a command line interface which is part of the OS to manage MPIO devices and the corresponding I/O paths. A Default PCM is available which provides a generic set of configuration methods to define, configure and remove MPIO devices. Also, this Default PCM provides a set of basic multi path functionalities like path detection, path failover, path health checking and load balancing.

The HP XP MPIIO solution for XP disk arrays utilizes the Default PCM of IBM and provides an AIX installable package which introduces additional ODM definitions in AIX. These ODM definitions are required to specify XP disk devices as MPIIO capable devices within AIX. Once the ODM definitions are in place AIX will detect and configure XP disk devices as MPIIO devices including any existing alternate paths. Standard MPIIO commands can be used to manage these devices.

The current MPIIO solution has the following behavioral constraints:

- The MPIIO solution cannot coexist with other non-MPIIO multipathing products on the same server. The MPIIO solution cannot coexist with HDLM on the same server.
- SCSI-3 persistent reservations are not supported.
- HACMP is supported when using enhanced concurrent volume groups within the cluster. Volume groups not controlled by the cluster can be standard volume groups.
- After re-enabling previously failed paths it can take a while (few minutes) until MPIIO takes all IO paths online again. This is not a defect.
- In the progress of an initial open to a raw MPIIO disk device MPIIO will check all IO paths and tries to re-enable previously failed paths automatically. This is also not considered to be a defect.

Check the SPOCK Website for the latest edition of this document. Go to <http://spock.corp.hp.com> and look under Quick Connect Guides.

Support

System requirements

This section describes the environment which is supported by the HP XP MPIO solution for XP disk arrays at the time of writing.

XP disk arrays and firmware version supported:

XP Disk Array	Firmware Version	Device Description
XP512/XP48	All Supported Versions	HP MPIO Disk (Fibre)
XP1024/XP128	All Supported Versions	HP MPIO Disk (Fibre)
XP12000/XP10000	All Supported Versions	HP MPIO Disk XP12000 (Fibre)
XP24000/XP20000	All Supported Versions	HP MPIO Disk XP24000 (Fibre)

Only IBM branded HBAs are supported. Please refer to SPOCK to see details on HBA support.

OS levels plus maintenance levels required for MPIO solution version 5.4.0.2

Early in 2006, IBM started moving away from the term “Maintenance Level” for the AIX OS, and newer versions of the AIX operating system are using the terms “Technology Level”, and a subset of Technology Level is called “Service Pack”. All necessary patches are included in the below listed Technology and Service Pack levels.

OS Level	Technology Level and Service Pack	Notes
AIX 5.2	TL08, SP02 or greater	For AIX 5.2, TL10, it is necessary to install SP02, or greater.
AIX 5.3	TL04, SP01 or greater	For AIX 5.3, TL06, it is necessary to install SP02, or greater.
AIX 6.1	TL00, SP01 or greater	AIX 6.1 is not supported on all XP disk arrays. Check the Streams documents for supported arrays.
P5 VIOS	ioslevel 1.2.1.1	Check with IBM for the currently supported versions of VIOS.

To verify the OS level plus maintenance level use the command **oslevel**. The following example shows an OS level of 5.3 with technology level 4, and Service Pack 2:

```
# oslevel -s  
5300-04-02
```

To check the existence of a specific patch on a system run the command **instfix**. The example shows that the OS patch “APAR IY42816” was successfully installed:

```
# instfix -ik IY42816
All filesets for IY42816 were found.
```

To verify the level of a specific fileset installed use the command **lspp**. This example shows that the level of the fileset `devices.common.IBM.mpio.rte` is 5.2.0.87:

```
# lspp -l devices.common.IBM.mpio.rte
Fileset                Level  State      Description
-----
Path: /usr/lib/objrepos
devices.common.IBM.mpio.rte
                        5.2.0.87  COMMITTED  MPIO Disk Path Control Module
```

SAN Boot

SAN boot is supported with the HP XP MPIO solution. See Streams documents for more detail.

Dynamic Tracking

Dynamic Tracking is supported with the HP XP MPIO solution.

Coexistence

Single Path ODM

For non-MPIO environments you must not install the HP XP MPIO solution. For single path environments and HDLM XP an additional ODM fileset is provided by HP. This fileset introduces a dedicated device type in the AIX operating system. This fileset provides additional ODM device definitions where the default values of the attributes are adjusted according to the XP disk array requirements. The Single Path ODM fileset cannot coexist with the fileset of the HP XP MPIO solution on the same server.

Nevertheless the HP XP MPIO solution can be used in single path or single HBA environments. Please be aware that in this case you have reduced or no path redundancy.

Check to see if the single path ODM is installed on an AIX system:

```
# lspp -l HP.aix.support.rte
```

If the fileset `HP.aix.support.rte` is installed on the system, it needs to be uninstalled before installing MPIO.

HDLM XP (Hitachi Dynamic Link Manager for XP arrays)

The HP XP MPIO solution cannot coexist with HDLM XP on the same server. Note that the Single Path ODM is required for HDLM.

Check to see if the HDLM software is installed on an AIX system:

For systems with HDLM version 5.70 or greater installed:

```
# lslpp -l DLManager.mpio.rte
```

For systems with HDLM version 5.64 or older installed:

```
# lslpp -l DLManager.rte
```

If any of the above HDLM filesets are installed on the system, The HDLM software needs to be uninstalled before installing MPIO.

Other MPIO implementations

The HP XP MPIO solution can coexist with the HP EVA MPIO solution on the same server utilizing the same HBA.

Check the SAN Design Guide for other supported coexistence configurations.

VIO Server

MPIO version 5.4.0.2 is supported on the VIO server. It is necessary to execute the `oem_setup_env` command. Please refer to IBM documentation concerning the usage of the `oem_setup_env` command.

Micro Partitioning

The HP XP MPIO solution supports the micro partitioning functionality introduced with the p5-Series of IBM on AIX 5.3

HACMP

When a shared volume group in an HACMP cluster is accessed through MPIO, it must be defined as an enhanced concurrent volume group. In an enhanced concurrent volume group environment no SCSI reservations are used to control access to the shared disk devices within the cluster. This eliminates the need to provide a utility to break SCSI reservations in the event of a node failover.

The standard LVM cannot be used for shared disk devices in an HACMP cluster utilizing the HP XP MPIO solution because there are no means to control SCSI reservations and potential ghost disks.

It is important to note that this does not limit the use of MPIO with HACMP to concurrent resource groups. Starting with HACMP V5.1, and continuing with HACMP 5.2, enhanced concurrent mode volume groups can be used with nonconcurrent resource groups. This facility, referred to as "Fast Disk Takeover" in the HACMP publications, combines the recovery speed of concurrent volume groups with the secure, single system access of non-concurrent volume groups. However, when this facility is used, it is recommended that the customer employ SAN zoning or physical connectivity schemes to ensure that only systems that are part of the HACMP cluster have access to the volume group(s).

Raid Manager

MPIO disk devices are supported as Raid Manager command devices. In this setup the failover mechanism is implemented by MPIO not Raid Manager.

MPIO Functionality

- The HP XP MPIO solution is compiled in a single fileset named “**devices.fcp.disk.HP.xparray.mpio.rte**”. Version 5.4.0.1 is an update to version 5.4.0.0 and must be installed on top of version 5.4.0.0, or both versions can be installed at the same time. Version 5.4.0.2 is also an update and can be installed with the prior versions at the same time.

The HP XP MPIO solution is based on the default PCM provided by IBM. This default PCM provides some basic multi path functionality. Because of certain limitations (e.g. no persistent reservations) not all functionalities of the default PCM are being utilized in this version. The following table gives an overview of the basic multi path functionality provided by the HP XP MPIO solution.

Multipath Functionality	Description
Automatic Discovery and Configuration	The MPIO architecture is integrated into the standard AIX device configuration methods. XP MPIO disk devices and all their associated io paths are configured automatically during system boot or when running <code>cfgmgr</code> . There is no need to serialize the configuration of the XP disk devices like in older HDLM environments (i.e. <code>cfgmgr -l fcs0</code> ; <code>cfgmgr -l fcs1</code>) because no SCSI reservation conflicts can occur with MPIO devices.
Automatic Failover	In the event of a failure MPIO will redirect all IOs to the alternate io path within a few seconds.
Automatic Failback	HP XP MPIO version 5.4.0.2 does provide load balancing. The default configuration of MPIO provides only single path access, therefore all IO will be sent via a single io path to the device. The primary path will always be the io path detected first. If the primary path fails MPIO will fail over to the next io path in the list (path id). Once the original path recovers, MPIO will not fail back to the original path automatically. Only during the initial open of the disk device, will MPIO utilize the primary path automatically again.
Load Balancing (round-robin)	Load balanced I/O is now supported. Set the attribute “algorithm” to round-robin in addition to setting “reserve_policy” to no_reserve for each hdisk in order for load balanced I/O to access the array Lun.
Path Health Checking	Path health checking is supported in versions 5.4.0.1 and 5.4.0.2.
SCSI Reservations	HP XP MPIO utilizes SCSI-2 reservations (single_path) which does not allow multiple initiators to access a device. Therefore the “reserve_policy” attribute of an hdisk must be set to no_reserve, in order for load balancing to work. SCSI-3 persistent reservations (PR_exclusive) are not supported.
Preferred Path	The HP XP MPIO solution does not provide a mechanism to specify a preferred path.

HP XP MPIO solution provides default values for the XP MPIO disk devices. These values are important for the correct operation and should not be modified unless specified. For a better understanding of the functionality provided by this package an extract of the attributes are explained below.

Attribute	Default Value	Description
PR_key_value	None	Key for persistent reservations. Since we are not supporting PR at this point in time no key is provided.
Algorithm	fail_over	The load balancing algorithm, “fail_over”, will send all IO over a single path. All other paths are standby paths. The load balancing algorithm “round_robin” will send I/O over all available paths. The attribute, reserve_policy, will need to be set to “no_reserve”.
hcheck_interval	60	The path health check interval is set to 60 seconds by default when the package is installed.
hcheck_mode	nonactive	This parameter specifies which io paths should be monitored by the path health checking functionality. nonactive: This mode will check all io paths of status failed and all standby paths of used/opened devices. Non-opened devices are not monitored. failed: This mode will check failed io paths only. enabled: This mode will check all enabled io paths which are opened. Non-opened devices are not monitored.
queue_depth	2	The queue depth can be adjusted according to the tuning requirements. Consider the service queue of the XP FC port and the maximum of 32 per LUN.
reserve_policy	Single_path	Standard SCSI-2 reservations are being used because persistent reservations are not supported at this point in time. This setting should be changed to “no_reserve” when the algorithm attribute is set to “round_robin”.
rw_timeout	60	The read/write timeout is 60 seconds. This setting should not be changed.

Installation

You will find the installation image of the HP XP MPIO solution on the XP firmware CD at `\TOOLS_PACK\CONFIG\ODM\AIX5L\ XPMPIODM.TAR`.

The file XPMPIO.TAR needs to be transferred onto the server you wish to install the HP XP MPIO solution. The file size of the tar archive is 60 KB in total. Once the file is located on the local server you can install the fileset via command line or via SMIT. Place the install package into a separate directory and create a TOC file. Then install the package using the standard AIX install utility `installp`.

Note that the installation of MPIO version 5.4.0.2 involves three filesets. They are MPIO 5400I, 5401U and 5402U. If MPIO 5401U is already installed on the system, then MPIO 5402U can be installed over it. The “U” means update. If 5400I and 5401U are not already installed on the system, then all three filesets can be installed at the same time. If any other versions of MPIO for XP are installed, they will have to be uninstalled before installing the 54xx versions of MPIO. A reboot is necessary after the installation.

Check for XP24000 devices already configured on the system

MPIO for XP version 5.4.0.1 displays hdisks configured for the XP24000 and XP20000 disk arrays as XP12000 devices when the command “`lsdev -Cc disk`” is used. While MPIO for XP version 5.4.0.2 can be installed over version 5.4.0.1, the device identifier string will not change - even after a reboot. In order to change the identifier string to the correct XP24000, the device will have to be removed from the system and `cfgmgr` run again. Note that there is no difference in the functionality between version 5.4.0.1 and 5.4.0.2. Version 5.4.0.2 now correctly identifies the XP24000 array.

```
# rmdev -dl hdiskX
# cfgmgr
```

Check for the non-MPIO XP ODM fileset

```
# lslpp -L HP.aix.support.rte
```

If this fileset exists, then it must be uninstalled.

Install via Command Line

```
# mkdir /tmp/hmpio # or any working directory
# cp XPMPIODM.TAR /tmp/hmpio # from CD
# cd /tmp/hmpio
# tar xvf XPMPIODM.TAR
# inutoc $PWD
# installp -acd . -einstallp.log ALL
```

Install via SMIT

```
# mkdir /tmp/hpmpio          # or any working directory
# cp XPMPIODM.TAR /tmp/hpmpio # from CD
# cd /tmp/hpmpio
# tar xvf XPMPIODM.TAR
# inutoc $PWD
# smitty install_latest
```

Install Software			
Type or select a value for the entry field. Press Enter AFTER making all desired changes.			
* INPUT device / directory for software	[.]	[Entry Fields]	+
F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Press Enter

Install Software			
Type or select values in entry fields. Press Enter AFTER making all desired changes.			
* INPUT device / directory for software	.	[Entry Fields]	
* SOFTWARE to install	[_all_latest]		+
PREVIEW only? (install operation will NOT occur)	no		+
COMMIT software updates?	yes		+
SAVE replaced files?	no		+
AUTOMATICALLY install requisite software?	yes		+
EXTEND file systems if space needed?	yes		+
OVERWRITE same or newer versions?	no		+
VERIFY install and check file sizes?	no		+
Include corresponding LANGUAGE filesets?	yes		+

---- end ----

Check Installation

Check the install log for any errors. If you installed via SMIT, check the file \$HOME/smit.log. If you installed via command line and used the option “-installp.log”, you need to check the file installp.log in the current directory.

Also you can check that status of the fileset with the command lslpp:

```
# lslpp -L devices.fcp.disk.HP.xparray.mpio.rte
Fileset                      Level  State  Type  Description (Uninstaller)
-----
devices.fcp.disk.HP.xparray.mpio.rte
                               5.4.0.2  C      F      AIX MPIO Support for XP Disk
                               Array (Update)
```

State codes:

- A -- Applied.
- B -- Broken.
- C -- Committed.
- E -- EFIX Locked.
- O -- Obsolete. (partially migrated to newer version)
- ? -- Inconsistent State...Run lppchk -v.

Type codes:

- F -- Installp Fileset
- P -- Product
- C -- Component
- T -- Feature
- R -- RPM Package

NOTE: A reboot will be necessary - read on.

Device Configuration

Delete old XP disk devices:

If the system has already discovered XP disk devices prior the installation of the HP XP MPIO solution you will need to remove these devices in order to reconfigure the disk devices as MPIO capable devices. The following example shows how to delete a single instance of a disk device.

```
# rmdev -dl hdisk10
```

This example shows how to delete all disks attached to a FC adapter fscsi0.

```
# rmdev -R -dl fscsi0
```

Customize FC Adapter Settings:

It is strongly recommended to enable the fast failover behavior of the FC SCSI driver (fscsix) by setting the device attribute **fc_err_recov** to **fast_fail**. This attribute needs to be adjusted manually since the default setting of this attribute is **delayed_fail**. If this parameter is not set correctly you may experience very long failover times in the event of a link failure. This parameter can only be changed if the device (fscsix) is not busy. Therefore all child devices (disks) need to be deleted before adjusting the device attribute. See above for instructions on deleting disk devices.

List all FC SCSI adapter drivers:

```
# lsdev -Ct efscsi
fscsi0 Available 1Z-08-02 FC SCSI I/O Controller Protocol Device
fscsi1 Available 1n-08-02 FC SCSI I/O Controller Protocol Device
```

List the current settings of all attributes of the FC SCSI adapter driver fscsi0:

```
# lsattr -El fscsi0
attach          switch          How this adapter is CONNECTED          False
dyntrk         no             Dynamic Tracking of FC Devices         True
fc_err_recov delayed_fail FC Fabric Event Error RECOVERY Policy True
scsi_id        0x7c0200      Adapter SCSI ID                        False
sw_fc_class    3             FC Class for Fabric                    True
```

To change the value of the attribute **fc_err_recov** to **fast_fail**:

```
# chdev -a fc_err_recov=fast_fail -l fscsi0
```

Discover XP MPIO disk devices:

After successfully installing the HP XP MPIO fileset(s), you must REBOOT the system. If a reboot has occurred since the install of the MPIO ODM, you may discover disk devices simply by calling the standard AIX device configuration utility **cfgmgr**. There is no need to serialize the scanning of the disk devices by scanning one HBA after the other in order to avoid any potential SCSI reservations because MPIO capable devices do not cause any SCSI reservation conflicts during the scan phase.

```
# cfgmgr
```

List XP MPIO disk devices and corresponding io paths:

```
# lsdev -Ct "xp*"
hdisk2 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
hdisk3 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
hdisk4 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
hdisk5 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
```

```
# lspath
Enabled hdisk2 fscsi0
Enabled hdisk3 fscsi0
Enabled hdisk4 fscsi0
```

```
Enabled hdisk5 fscsi0
Enabled hdisk2 fscsi1
Enabled hdisk3 fscsi1
Enabled hdisk4 fscsi1
Enabled hdisk5 fscsi1
```

```
# lspath -F "name parent path_id status path_status connection" -l hdisk3
hdisk3 fscsi0 0 Enabled Available 50060e8004276964,10000000000000
hdisk3 fscsi1 1 Enabled Available 50060e8004276974,10000000000000
```

```
# lsattr -El hdisk2
```

PCM	PCM/friend/xparray	Path Control Module	False
PR_key_value	none	Reserve Key	True
algorithm	fail_over	Algorithm	True
clr_q	no	Device CLEARS its Queue	on
error	True		
hcheck_interval	60	Health Check Interval	True
hcheck_mode	nonactive	Health Check Mode	True
location		Location Label	True
lun_id	0x0	Logical Unit Number ID	False
lun_reset_spt	yes	SCSI LUN reset	True
max_transfer	0x40000	N/A	True
node_name	0x50060e8004276964	Node Name	False
pvid	0052198f5c8706740000000000000000	Physical Volume ID	False
q_err	yes	Use QERR bit	False
q_type	simple	Queue TYPE	True
queue_depth	2	Queue DEPTH	True
reassign_to	120	REASSIGN time out	True
reserve_policy	single_path	Reserve Policy	True
rw_timeout	60	READ/WRITE time out	True
scsi_id	0x790000	SCSI ID	False
start_timeout	60	START UNIT time out	True
ww_name	0x50060e8004276964	FC World Wide Name	False

Customize XP MPIO disk devices:

To allow round robin load balancing over multiple paths, make the following changes, per Lun:

```
# chdev -l hdisk2 -a reserve_policy=no_reserve
hdisk2 changed
```

```
# chdev -l hdisk2 -a algorithm=round_robin
hdisk2 changed
```

Adjust the queue depth according to your needs:

```
# chdev -l hdisk2 -a queue_depth=8
hdisk2 changed
```

and the check:

```
# lsattr -El hdisk2 -a queue_depth
queue_depth 8 Queue DEPTH True
```

Software Removal

Before you remove the fileset of the HP XP MPIO solution you must remove all the XP MPIO disk devices from the system. Otherwise you will end up with orphan disk device definitions which cannot be managed by AIX. To be able to remove a disk device from the system the device must not be opened. Therefore deactivate all corresponding file systems and volume groups and stop all corresponding raid manager instances.

To remove a single disk device type:

```
# rmdev -dl hdisk2
hdisk2 deleted
```

To delete all HP XP MPIO disk devices type:

```
# lsdev -Ct "xp*" -F name | xargs -l rmdev -dl
```

After successful removal of all HP XP MPIO disk devices the fileset of the HP XP MPIO solution can be removed via command line

```
# installp -u devices.fcp.disk.HP.xparray.mpio.rte
```

Or via SMIT:

```
# smit "maintain_software"
```

```
Software Maintenance and Utilities

Move cursor to desired item and press Enter.

Commit Applied Software Updates (Remove Saved Files)
Reject Applied Software Updates (Use Previous Version)
Remove Installed Software

Rename Software Images in Repository
Clean Up Software Images in Repository

Copy Software to Hard Disk for Future Installation
Copy Software Bundle to Hard Disk for Future Installation

Check Software File Sizes After Installation
Verify Software Installation and Requisites

Clean Up After Failed or Interrupted Installation

F1=Help           F2=Refresh        F3=Cancel         F8=Image
F9=Shell          F10=Exit          Enter=Do
```

Select "Remove Installed Software".

Remove Installed Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* SOFTWARE name	[]	+
PREVIEW only? (remove operation will NOT occur)	no	+
REMOVE dependent software?	no	+
EXTEND file systems if space needed?	no	+
DETAILED output?	no	+

F1=Help	F2=Refresh	F3=Cancel	F4=List
F5=Reset	F6=Command	F7=Edit	F8=Image
F9=Shell	F10=Exit	Enter=Do	

Press F4 on "SOFTWARE name" to get a list of all software installed. Use the keys "/" and "n" to search for the fileset devices.fcp.disk.HP.xparray.mpio.rte. Select the fileset with F7. Set "PREVIEW only?" to "no" and press enter.

MPIO Operation

There are two MPIO specific OS commands required to manage disk devices

1. [lspath](#) – lists the status and attributes of MPIO disk devices
2. [chpath](#) – enables or disables I/O paths or changes I/O path specific attributes
3. [rmpath](#) – deletes an I/O path

Some additional OS commands are required to manage MPIO disk devices

1. [cfgmgr](#) – configures new devices into the system (ioscan on HP-UX)
2. [lspv](#) – lists all physical volumes with PVID and volume group affinity
3. [lsdev](#) – lists devices
4. [lscfg](#) – lists configuration and VPD of a device
5. [lsattr](#) – lists the attributes of a device
6. [chdev](#) – changes the attributes of a device
7. [rmdev](#) – removes a device from the system
8. [iostat](#) – lists io statistics

Examples

lspath

1. To list all paths on the system

```
l# lspath
Enabled hdisk0 scsi0
Enabled hdisk1 scsi0
Enabled hdisk2 fscsi0
Enabled hdisk3 fscsi0
Enabled hdisk4 fscsi0
Enabled hdisk2 fscsi1
Enabled hdisk3 fscsi1
Enabled hdisk4 fscsi1
```

2. To list all the paths of a single disk device

```
# lspath -l hdisk4
Enabled hdisk4 fscsi0
Enabled hdisk4 fscsi1
```

3. To list all paths with additional attributes

```
# lspath -F "name parent path_id status path_status connection"
hdisk0 scsi0 0 Enabled Available 8,0
hdisk1 scsi0 0 Enabled Available 9,0
hdisk2 fscsi0 0 Enabled Available 50060e8004276964,1000000000000
hdisk3 fscsi0 0 Enabled Available 50060e8004276964,2000000000000
hdisk4 fscsi0 0 Enabled Available 50060e8004276964,3000000000000
hdisk2 fscsi1 1 Enabled Available 50060e8004276974,1000000000000
hdisk3 fscsi1 1 Enabled Available 50060e8004276974,2000000000000
hdisk4 fscsi1 1 Enabled Available 50060e8004276974,3000000000000
```

4. To list all paths of a single HBA

```
# lspath -p fscsi1
Enabled hdisk2 fscsi1
Enabled hdisk3 fscsi1
Enabled hdisk4 fscsi1
```

1. To list the attributes of a specific path of a disk device

```
# lspath -EAl hdisk4 -p fscsi0
scsi_id    0x790000          SCSI ID    False
node_name  0x50060e8004276964 Node Name  False
priority   1                Priority   True
```

chpath

2. To disable an I/O path of a disk device

```
# chpath -l hdisk4 -p fscsi1 -s disable
paths Disabled
```

3. To re-enable all I/O paths of a disk device

```
# chpath -l hdisk4 -s enable
lpaths Enabled
```

4. To re-enable a specific I/O path of a disk device

```
# chpath -l hdisk4 -p fscsi1 -s enable
paths Enabled
```

5. To disable all I/O paths of a specific adapter (fscsi0)

```
# lspath -p fscsi0 -F name | xargs -l chpath -s disabled -p fscsi0 -l
paths Disabled
paths Disabled
paths Disabled
```

rmpath

1. To delete an I/O path of a specific disk device

```
# rmpath -l hdisk4 -p fscsi0 -d
path deleted
```

2. To delete all I/O paths configured via a specific FC-SCSI protocol adapter

```
# rmpath -p fscsi0 -d
paths deleted
```

cfgmgr

1. To scan for new devices (including disk devices)

```
# cfgmgr
```

lspv

1. To list all available physical volumes

```
# lspv
hdisk0          0052198f96eed673          rootvg          active
hdisk1          0052198fc6aa06a1          None            active
hdisk2          none                            None            active
hdisk3          0052198f157be268          vgxp01         active
hdisk4          0052198f157be58c          vgxp01         active
```

lsdev

1. To list all disk devices

```
# lsdev -Cc disk
hdisk0 Available 1S-08-00-8,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 1S-08-00-9,0 16 Bit LVD SCSI Disk Drive
hdisk2 Available 1Z-08-02      HP MPIO Disk XP12000 (Fibre)
hdisk3 Available 1Z-08-02      HP MPIO Disk XP12000 (Fibre)
hdisk4 Available 1Z-08-02      HP MPIO Disk XP12000 (Fibre)
```

2. To list all XP disk devices

```
# lsdev -Cc disk -F "name status location class subclass type description" -t
"xp*"
hdisk2 Available 1Z-08-02 disk fcp xp12kmpio HP MPIO Disk XP12000 (Fibre)
hdisk3 Available 1Z-08-02 disk fcp xp12kmpio HP MPIO Disk XP12000 (Fibre)
hdisk4 Available 1Z-08-02 disk fcp xp12kmpio HP MPIO Disk XP12000 (Fibre)
```

3. To list all FC disk devices

```
# lsdev -Cc disk -s fcp
hdisk2 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
hdisk3 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
hdisk4 Available 1Z-08-02 HP MPIO Disk XP12000 (Fibre)
```

lscfg

1. To list all devices

```
# lscfg
INSTALLED RESOURCE LIST

The following resources are installed on the machine.
+/- = Added or deleted from Resource List.
*   = Diagnostic support not available.

Model Architecture: chrp
Model Implementation: Multiple Processor, PCI bus

+ sys0                                System Object
+ sysplanar0                          System Planar
* pci4                                U0.2-P1.1      PCI Bus
* pci3                                U0.2-P1.1      PCI Bus
* pci2                                U0.1-P1        PCI Bus
* pci14                               U0.1-P2        PCI Bus
```

```

* pci15          U0.1-P2          PCI Bus
* pci1           U0.1-P1          PCI Bus
* pci11          U0.1-P2          PCI Bus
+ fcs0           U0.1-P2-I4/Q1      FC Adapter
* fcnet0         U0.1-P2-I4/Q1      Fibre Channel Network
Protocol Device
* fscsi0         U0.1-P2-I4/Q1      FC SCSI I/O Controller
Protocol Device
* hdisk2         U0.1-P2-I4/Q1-W50060E8004276964-L1000000000000 HP MPIO Disk XP12000 (Fibre)
* hdisk3         U0.1-P2-I4/Q1-W50060E8004276964-L2000000000000 HP MPIO Disk XP12000 (Fibre)
* hdisk4         U0.1-P2-I4/Q1-W50060E8004276964-L3000000000000 HP MPIO Disk XP12000 (Fibre)
* pci13          U0.1-P2          PCI Bus
+ fcs1           U0.1-P2-I3/Q1      FC Adapter
...

```

2. To list all disk devices

```

# lscfg -l "hdisk*"
hdisk2 U0.1-P2-I4/Q1-W50060E8004276964-L1000000000000 HP MPIO Disk XP12000 (Fibre)
hdisk3 U0.1-P2-I4/Q1-W50060E8004276964-L2000000000000 HP MPIO Disk XP12000 (Fibre)
hdisk4 U0.1-P2-I4/Q1-W50060E8004276964-L3000000000000 HP MPIO Disk XP12000 (Fibre)
hdisk0 U0.1-P2/Z1-A8          16 Bit LVD SCSI Disk Drive (36400 MB)
hdisk1 U0.1-P2/Z1-A9          16 Bit LVD SCSI Disk Drive (36400 MB)

```

3. To list vital product data of a specific device

```

# lscfg -vl hdisk4
hdisk4 U0.1-P2-I4/Q1-W50060E8004276964-L3000000000000 HP MPIO Disk XP12000
(Fibre)

Manufacturer.....HP
Machine Type and Model.....OPEN-3
Part Number.....
ROS Level and ID.....35303034
Serial Number.....50 02769
EC Level.....
FRU Number.....
Device Specific.(Z0).....00000332CF000002
Device Specific.(Z1).....0162 7E ....
Device Specific.(Z2).....
Device Specific.(Z3).....
Device Specific.(Z4).....:1z
Device Specific.(Z5).....
Device Specific.(Z6).....

```

lsattr

1. To list attributes of a disk device

```

# lsattr -El hdisk4
PCM                PCM/friend/xparray      Path Control Module      False
PR_key_value       none                    Reserve Key               True
algorithm           fail_over              Algorithm                 True
clr_q              no                      Device CLEARS its Queue on error True
hcheck_interval    60                     Health Check Interval    True
hcheck_mode        nonactive              Health Check Mode        True
location            Location Label         Location Label            True
lun_id              0x3000000000000000    Logical Unit Number ID   False
lun_reset_spt      yes                    SCSI LUN reset           True
max_transfer        0x40000                N/A                       True
node_name           0x50060e8004276964    Node Name                 False
pvid                0052198f157be58c00000000000000000 Physical Volume ID       False
q_err              yes                    Use QERR bit             False
q_type             simple                 Queue TYPE                True

```

queue_depth	2	Queue DEPTH	True
reassign_to	120	REASSIGN time out	True
reserve_policy	single_path	Reserve Policy	True
rw_timeout	60	READ/WRITE time out	True
scsi_id	0x790000	SCSI ID	False
start_timeout	60	START UNIT time out	True
ww_name	0x50060e8004276964	FC World Wide Name	False

2. To list attributes of a FC-SCSI protocol adapter

```
# lsattr -El fscsi0
attach          switch      How this adapter is CONNECTED      False
dyntrk         no         Dynamic Tracking of FC Devices      True
fc_err_recov   fast_fail  FC Fabric Event Error RECOVERY Policy True
scsi_id        0x7c0200  Adapter SCSI ID                    False
sw_fc_class    3         FC Class for Fabric                True
```

3. To list the value of the attribute queue_depth of a disk device.

```
# lsattr -El hdisk4 -a queue_depth
queue_depth 2 Queue DEPTH True
```

chdev

1. To set the health check interval to 0 in order to disable MPIO path health checking. This task requires the device to be offline.

```
# umount /xyz
# varyoffvg vgxp01
# chdev -a hcheck_interval=0 -l hdisk4
hdisk4 changed
```

2. To change the queue depth of a disk device. This task requires the device to be offline.

```
# chdev -a queue_depth=4 -l hdisk4
hdisk4 changed
```

3. To change the FC error recovery from delayed failed to fast fail of the FC-SCSI protocol adapter (fscsi#). This task requires the device to be offline. Therefore all child devices configured on this adapter need to be deleted. Note, the alternate I/O path is still active.

```
# rmpath -p fscsi0 -d
paths deleted
# lsattr -El fscsi0 -a fc_err_recov
fc_err_recov delayed_fail FC Fabric Event Error RECOVERY Policy True
# chdev -a fc_err_recov=fast_fail -l fscsi0
fscsi0 changed
# cfgmgr -l fscsi0
# lsattr -El fscsi0 -a fc_err_recov
fc_err_recov fast_fail FC Fabric Event Error RECOVERY Policy True
```

rmdev

1. To remove a single disk device with all I/O paths. This task requires the device to be offline

```
# umount /xyz
# varyoffvg vgxp01
# rmdev -dl hdisk4
```

- ```

hdisk4 deleted
2. To remove all XP disk devices
lsdev -Ct "xp*" -F name | xargs -l rmdev -dl
hdisk2 deleted
hdisk3 deleted
hdisk4 deleted
3. To remove all disk devices configured via FC-SCSI protocol adapter fscsi1 plus FC-SCSI
protocol adapter itself
rmdev -R -dl fscsi0
fscsi0 deleted

```

## iostat

- To list disk I/O statistics

```
iostat -Td hdisk3 hdisk4 5 3
```

```
System configuration: lcpu=1 disk=5
```

| Disks: | % tm_act | Kbps    | tps   | Kb_read | Kb_wrtn  | time     |
|--------|----------|---------|-------|---------|----------|----------|
| hdisk4 | 1.3      | 1397.0  | 11.1  | 2110    | 10431624 | 12:46:25 |
| hdisk3 | 1.4      | 1401.9  | 11.2  | 6576    | 10463632 | 12:46:25 |
| hdisk4 | 37.2     | 39244.8 | 306.8 | 0       | 196224   | 12:46:30 |
| hdisk3 | 36.6     | 39321.6 | 307.2 | 0       | 196608   | 12:46:30 |
| hdisk4 | 36.0     | 37120.0 | 289.8 | 0       | 185600   | 12:46:35 |
| hdisk3 | 35.4     | 37683.2 | 294.4 | 0       | 188416   | 12:46:35 |

- To list disk I/O statistics with detailed path throughput. This enables you to check for the active I/O path. In this example Path1 is the active path for hdisk3 and hdisk4.

```
iostat -Tdm hdisk3 hdisk4 5
```

```
System configuration: lcpu=1 disk=5
```

| Disks: | % tm_act | Kbps  | tps | Kb_read | Kb_wrtn | time     |
|--------|----------|-------|-----|---------|---------|----------|
| hdisk4 | 0.7      | 662.4 | 5.4 | 2110    | 4838920 | 12:43:45 |

| Paths: | % tm_act | Kbps  | tps | Kb_read | Kb_wrtn | time     |
|--------|----------|-------|-----|---------|---------|----------|
| Path0  | 0.0      | 0.0   | 0.0 | 0       | 0       | 12:43:45 |
| Path1  | 0.7      | 662.4 | 5.4 | 2110    | 4838920 | 12:43:45 |

| Disks: | % tm_act | Kbps  | tps | Kb_read | Kb_wrtn | time     |
|--------|----------|-------|-----|---------|---------|----------|
| hdisk3 | 0.6      | 665.3 | 5.4 | 6576    | 4855568 | 12:43:45 |

| Paths: | % tm_act | Kbps  | tps | Kb_read | Kb_wrtn | time     |
|--------|----------|-------|-----|---------|---------|----------|
| Path0  | 0.0      | 0.0   | 0.0 | 0       | 0       | 12:43:45 |
| Path1  | 0.6      | 665.3 | 5.4 | 6576    | 4855568 | 12:43:45 |

| Disks: | % tm_act | Kbps    | tps   | Kb_read | Kb_wrtn | time     |
|--------|----------|---------|-------|---------|---------|----------|
| hdisk4 | 30.2     | 35712.0 | 279.0 | 0       | 178560  | 12:43:50 |

| Paths: | % tm_act | Kbps    | tps   | Kb_read | Kb_wrtn | time     |
|--------|----------|---------|-------|---------|---------|----------|
| Path0  | 0.0      | 0.0     | 0.0   | 0       | 0       | 12:43:50 |
| Path1  | 30.2     | 35712.0 | 279.0 | 0       | 178560  | 12:43:50 |

|        |          |         |       |         |         |          |
|--------|----------|---------|-------|---------|---------|----------|
| Disks: | % tm_act | Kbps    | tps   | Kb_read | Kb_wrtn | time     |
| hdisk3 | 43.8     | 35404.8 | 276.6 | 0       | 177024  | 12:43:50 |
| Paths: | % tm_act | Kbps    | tps   | Kb_read | Kb_wrtn |          |
| Path0  | 0.0      | 0.0     | 0.0   | 0       | 0       | 12:43:50 |
| Path1  | 43.8     | 35404.8 | 276.6 | 0       | 177024  | 12:43:50 |

## Release History

---

The HP XP MPIO solution, version 5.4.0.2, has the following enhancements:

- The MPIO software now correctly identifies the XP24000 disk array.

The HP XP MPIO solution, version 5.4.0.1, has the following enhancements:

- Path health checking is now available and the check interval is user settable.
- MPIO load balancing is supported in a round-robin fashion. Reservations must be turned off.
- The MPIO solution can coexist with other MPIO packages on the same server. This allows connecting to different storage device types.

The HP XP MPIO version 5.4.0.0 has the following constraints:

- The Path Health Checking feature is unsupported and must be turned off after installation.
- Failed paths, once repaired, must be re-enabled by executing a “chpath” command from the command line.
- Load balanced I/O is not available. I/O can only be issued over a single path to the array Lun.