



Readme:
3PAR[®] Multipath I/O 2.2 for IBM AIX

For Printed Material

Copyright © 2006 3PARdata, Inc. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of 3PARdata, Inc., 4209 Technology Drive, Fremont, CA 94538.

For Electronic Material

The user may print one copy for personal use. Production of multiple copies or copies for sale is strictly prohibited.

Trademarks

3PARdata, the 3PAR logo, 3PAR, InServ, InForm, and InSpire are trademarks of 3PARdata, Inc. All other trademarks and registered trademarks are the property of their respective owners.

Introduction

3PAR® MPIO for IBM AIX enables the use of more than one physical I/O path to access a 3PAR InServ Storage Server. By offering fault tolerance and/or load balancing of I/O traffic with Active/Active support, multipathing provides improved system reliability and availability.

The contents included within the CD are:

readme.doc	This file.
3PARmpio.bff	3PARmpio.64 package.
.toc	Installation table of content.

Supported IBM AIX Versions

The current 3PAR MPIO for IBM AIX supports the following AIX versions:

<u>AIX version</u>	<u>oslevel -r</u>
AIX 5.2 ML6	5200-06
AIX 5.3 ML3	5300-03

Refer to the 3PAR Configuration Matrix and the 3PAR Best Practice documents for IBM AIX for complete details on supported configurations.

Installation

Perform the installation of 3PAR MPIO for IBM AIX via the `smit install`.

For hosts containing 3PAR MPIO 2.1 for IBM AIX, the installation will handle upgrading from 3PAR MPIO 2.1 to 3PAR MPIO 2.2 for IBM AIX automatically. There is no need to uninstall 3PAR MPIO 2.1 for IBM AIX.

For hosts containing 3PAR AIX ODM 1.1, the installation will handle upgrading from 3PAR AIX ODM 1.1 to 3PAR MPIO 2.2 for IBM AIX automatically. There is no need to uninstall 3PAR AIX ODM 1.1.

For hosts containing 3PAR AIX ODM 1.0 or 1.2, deinstall the software prior to installing the 3PAR MPIO 2.2 for IBM AIX. The package contains a pre-installation check and the installation will fail if 3PAR AIX ODM 1.0 or 1.2 version is found.

NOTE: Installation of this package requires a reboot of the host AIX server.

To install the 3PAR MPIO for IBM AIX:

1. Insert the installation CD into the CD-ROM drive.
2. Execute the `smit install`. Select **Install and Update Software > Install Software**. Press **F4** to select `source /dev/cdX`, where `X` is the device number. For example, `/dev/cd0`. Click **Enter**. The `smit` will then install the 3PAR MPIO software. Upon completion, the command status displays `OK`.
3. For AIX 5.2 ML1 and later, IBM introduced the dynamic tracking features. The dynamic tracking defaults to disabled by AIX. The installation adds the required 3PAR ODM attributes for the dynamic tracking; however, it will not enable it. To enable the dynamic tracking, execute the `/usr/lpp/3PARmpio/bin/3par_dyntrk.sh` script.

NOTE: For an AIX host that will be attached to the 3PAR InServ Storage Server for the first time, the `/usr/lpp/3PARmpio/bin/3par_dyntrk.sh` script will not find any 3PAR device. Therefore no dynamic tracking will be enabled on any adapter. Follow through the rest of the steps. After reboot, the 3PAR devices are configured and recognized by AIX host. Run the `3par_dyntrk.sh` script at this point. It enables the dynamic tracking on the adapters that have 3PAR devices attached. Then reboot. The dynamic tracking features will be active after the reboot. The `3par_dyntrk.sh` script also comes with the `-a` option, which enables the dynamic tracking features on all adapters.

4. In order to see the 3PAR device(s) from the AIX host, proper port persona on the InServ Storage Server should be set. Refer to appropriate 3PAR best practice document.
5. Reboot the system.

Post-Installation

To verify the 3PAR MPIO package after reboot, please issue `lspp -l 3PARmpio.64`. This command displays the package level and state. To list the 3PAR MPIO devices, please issue `lsdev -Cc disk`. The 3PAR device is described as `3PAR InServ Virtual Volume`.

De-Installation

To deinstall the 3PAR MPIO package for IBM AIX:

1. Remove all the devices based on 3PAR VLUNs. For example, `rmdev -dl hdiskX`, where `hdiskX` is a 3PAR InServ Storage Server virtual volume. The following is a sample script to remove all 3PAR devices:

```
# for i in $(lsdev -Cc disk |awk '/3PAR InServ Virtual Volume/{print $1}')
# do
# rmdev -dl $i
# done
```

2. Execute the `smit install`. Select **Software Maintenance and Utilities > Remove Installed Software**. Press **F4** to select the `3PARmpio.64`. Change the **PREVIEW only?** to **no**. Click **Enter**. The `smit` will then deinstall the 3PAR MPIO software. Upon completion, the command status displays `OK`.
3. Reboot the system.

Host Boot from an InServ Storage Server

NOTE: Once 3PAR MPIO 2.2 for IBM AIX has been installed, it cannot be de-installed for a host boot disk.

AIX host boot is supported by the InServ Storage Server. However, in order to successfully install the AIX operating system on InServ volumes, the following steps must be followed:

1. Configure the InServ Storage Server to the AIX system. Refer to the 3PAR Best Practice documents for IBM AIX for detailed instructions.

NOTE: Restrict a single path connection between the InServ Storage Server and the AIX host system.

2. Install the AIX operating system on the selected InServ volume.
3. Reboot the AIX system. The InServ volume is configured as an MPIO capable device with the AIX default PCM. Upgrade the AIX operating system to the recommended maintenance level. Refer to the 3PAR Best Practice documents for IBM AIX for detailed instructions.
4. Install the 3PAR MPIO 2.2 for IBM AIX.
5. Execute `bosboot -aDd /dev/ipldevice`.
6. Reboot the AIX system.
7. After the AIX system comes up completely, connect additional paths to the Fabric or the 3PAR disk storage system. Configure additional paths.

Examples of configuring additional paths:

```
/usr/lib/methods/cfgefscsi -l fscsiX
```

In the example above, `x` is the additional path's FC SCSI controller protocol device.

```
cfgmgr -vl hdiskX
```

In the example above, `x` is the InServ volume for the host boot.

Example of displaying multiple paths:

```
lspath -l hdiskX
```

In example above, X is the InServ volume for the host boot

8. Execute `bosboot -aDd /dev/ipldevice`.
9. Reboot the AIX system.

All InServ volumes, including the selected InServ boot volumes, are now configured with the 3PAR PCM.

Utility Scripts

There are some useful 3PAR utility scripts that reside in the `/usr/lpp/3PARmpio/bin`:

`3par_dyntrk.sh [-a]` - This script enables the dynamic tracking features on adapters that have 3PAR devices attached. If it is preferable to enable dynamic tracking on all adapters, the `-a` option can be used.

`3par_hdisk2vv.sh` - This script takes an AIX hdisk as input and converts it to its associated 3PAR virtual volume.

`3par_vv2hdisk.sh` - This script takes a 3PAR Virtual Volume (VV) as input and converts it to its associated AIX hdisk.

`3par_pathmon.sh [-d -s -r -a] [-b] [-i interval]:` - This daemon-like program is designed to monitor the 3PAR device paths in Failed state. For all Failed state paths, which the health check function cannot bring back, a `chpath` is issued to enable these paths. By default the program will check the path state every 20 seconds. To change the interval, the program needs to be stopped by using the `-d` option, then restarted by using the `-i` option with the preferred time interval in seconds. The `-s` option checks the program status. The program has an entry named `3par_pathmon` in the `/etc/inittab` with once setting. To remove the entry, issue the `3par_pathmon.sh` script with the `-r` option. If the entry is removed by the `-r` option, the path monitoring program will NOT auto run in the next reboot. To re-add the entry, issue the `3par_pathmon.sh` script with the `-a` option. The `-b` option runs the program in debug mode. For details of the health check function, see the following section, *Manage 3PAR MPIO Device by SMIT*.

`3par_explorer.sh` - This script takes a system snapshot and compresses the output for debugging purpose.

CAUTION: The following two scripts are needed ONLY during the change of the AIX OS kernel mode. Incorrectly executing these scripts can cause the 3PAR device(s) to be un-configured.

`3par_enable32.sh` - Enables the 3PARmpio 32 bit package.

`3par_enable64.sh` - Enables the 3PARmpio 64 bit package.

Manage 3PAR MPIO Device by SMIT

The 3PAR AIX MPIO device can be managed by `smit mpio`. There are three categories under the `smit mpio` management:

- MPIO Path Management

- MPIO Device Management

- MPIO Parent Management

The common use for the Path Management is to enable the Failed status path. To enable the path, select **MPIO Path Management > Enable Paths > Enable Selected or All Paths**.

The Device Management can be used to change the default device attributes setting. The following are the default package settings for the 3PAR device:

Device Attribute	Attribute Value	Description
reserve_policy	no_reserve	Does not apply a reservation methodology for the device.
algorithm	round_robin	The default round_robin algorithm cannot use the single_reserve policy. Any attempt to set in this combination will fail.
hcheck_interval	10	<p>The health check interval is set to 10 seconds by default. It is the interval value that the MPIO device driver would send a SCSI command to the associated device. For example, if one of the device paths is recovered from the Failed state, the health check SCSI command will have a good return code and enable the path automatically. To disable this feature, change the <code>hcheck_interval</code> value to 0.</p> <p>NOTE: Currently the AIX MPIO healthcheck function only checks the paths being opened. It does not healthcheck any path that is in the close state. The MPIO device driver will change the path to the close state after I/O completes. As a result, if a path fails during the I/O and does not recover, the path will be put in the failed and close state after the I/O completes. The healthcheck will not recover the path in this case. However, the path can be manually enabled by using <code>smit mpio</code>. Select MPIO Path Management > Enable Paths > Enable Selected or All Paths.</p>
priority	1	All path priority is 1. It means that each path of the same device has the same priority when using the round_robin algorithm. It is the only priority value in this version.

max_transfer	0x40000	Maximum transfer size in bytes that can be transferred to the device in one operation. The default is 256K (0x40000).
--------------	---------	---

To change these attributes, select **MPIO Device Management > Change/Show MPIO Device Characteristics**. Select the Device Name by pressing **F4**. Change the attribute(s) value and then click **Enter**.

The followings are sample SMIT actions:

SMIT Action	Steps
Enable path	<ol style="list-style-type: none"> 1. Execute <code>smit mpio</code>. 2. Select MPIO Path Management > Enable Paths > Enable Selected or All Paths. 3. Press Enter. <hr/> <p>NOTE: If a path is in the "Missing" or "Defined" state, that path cannot be enabled and must be defined and configured before it can be used. For example, if a cable is unplugged or failed on a configured path and the host is rebooted, the path will become "Missing" after the reboot. In order to define and configure a "Missing" path, the root cause must be resolved first; in this case, the unplugged or failed cable. For additional information on path configuration, see the examples provided in the <i>Host Boot from an InServ Storage Server</i> section on page 3.</p>
Change device algorithm	<ol style="list-style-type: none"> 1. Execute <code>smit mpio</code>. 2. Select MPIO Device Management > Change/Show MPIO Device Characteristics. 3. Select the Device Name by pressing F4. 4. Move the cursor to the Algorithm attribute. 5. Press F4 to select preferred algorithm. 6. Press Enter.

Change device reserve_policy	<ol style="list-style-type: none"> 1. Execute <code>smit mpio</code>. 2. Select MPIO Device Management > Change/Show MPIO Device Characteristics. 3. Select the Device Name by pressing F4. 4. Move the cursor to the Reservation Policy attribute. 5. Press F4 to select preferred reservation policy 6. Press Enter.
Change device hcheck_interval	<ol style="list-style-type: none"> 1. Execute <code>smit mpio</code>. 2. Select MPIO Device Management > Change/Show MPIO Device Characteristics. 3. Select the Device Name by pressing F4. 4. Move the cursor to the HealthCheck Interval 0-3600 sec attribute. 5. Type in the preferred interval in seconds. 6. Press Enter.
Change device max_transfer	<ol style="list-style-type: none"> 1. Execute <code>smit mpio</code>. 2. Select MPIO Device Management > Change/Show MPIO Device Characteristics. 3. Select the Device Name by pressing F4. 4. Move the cursor to the Max Transfer Size attribute. 5. Press F4 to select the preferred value. 6. Press Enter. <hr/> <p>NOTE: The device must be in Defined state in order to change this attribute.</p>

NOTE: To learn more about SMIT, issue `smit` and select **Using SMIT (information only)** from the SMIT main menu.

Understanding 3PAR Support for HACMP

You can run 3PAR MPIO 2.2 for IBM AIX in concurrent and non-concurrent multi-host environments. The same LUNs on the 3PAR storage device can be attached to the AIX hosts managed by HACMP. The 3PAR MPIO 2.2 can support HACMP version 5.3 running on pSeries servers with AIX 5.3 ML3. Refer to the 3PAR Configuration Matrix and the 3PAR Best Practice documents for IBM AIX for complete details on supported configurations.

3PAR MPIO 2.2 supports SCSI-3 persistent reserve attributes. It can be applied on each hdisk where HACMP non-concurrent volume groups will be created. To check the persistent reserve key value that HACMP will use, issue the command:

```
odmget -q "uniquetype=disk/fcp/3PAR_VV_MPIO AND attribute=PR_key_value" PdAt
```

The output should look similar to the following:

```
PdAt :

uniquetype = "disk/fcp/3PAR_VV_MPIO"
attribute = "PR_key_value"
deflt = "6075429844412053580"
values = ""
width = ""
type = "R"
generic = "DU"
rep = "s"
nls_index = 87
```

To enable the SCSI-3 persistent reserve on hdisk, issue:

```
`chdev -l hdiskX -a reserve_policy=PR_exclusive`
```

where x is the device number such as hdisk3. You can also change the hdisk device Reservation Policy to `PR_exclusive` through `smit` as described in the section, *Manage 3PAR MPIO Device by SMIT*. If the HACMP non-concurrent volume groups reside on different hdisks, you need to change the `reserve_policy` attribute to `PR_exclusive` on all these different hdisks. After completing the `reserve_policy` attribute setting, the SCSI-3 persistent reserve will be applied on these hdisks when HACMP vary on the non-concurrent volume groups.

NOTE: The 3PAR MPIO 2.2 for IBM AIX default `reserve_policy` is set to `no_reserve`. In order to apply SCSI-3 persistent reserve on an HACMP non-concurrent volume group, you must follow the steps described above to change the device `reserve_policy` accordingly.

To disable the SCSI-3 persistent reserve, issue:

```
chdev -l hdiskX -a reserve_policy=no_reserve
```

where *x* is the device number such as `hdisk3`. You can also change the `hdisk` device Reservation Policy to `no_reserve` through `smit` as described in the section, *Manage 3PAR MPIO Device by SMIT*.

For HACMP configuration, refer to the *HACMP 5.3 Concepts and Facilities Guide*, or visit the HACMP 5.3 documentation website:

http://www-03.ibm.com/servers/eserver/pseries/library/hacmp_docs.html

NOTE: The APAR IY82920 is required for AIX 5.3 ML3 to fully support the HACMP fabric configuration.
