ATTO Technology, Inc.

ATTO Celerity MultiPath Director™ Installation and Operation Manual

ATTO Technology, Inc.
155 CrossPoint Parkway
Amherst, New York 14068 USA
www.attotech.com

Tel (716) 691-1999
Fax (716) 691-9353

Sales support: sls@attotech.com
Technical support: Monday -- Friday, 8am-6pm EST
techsupp@attotech.com
(716) 691-1999 ext. 242

02/2013 PRMA-0410-000

Contents

ATTO MultiPath Director™ is available for Windows, Linux and Mac OS X Operating Systems. Please consult your storage vendor for a complete list of supported ATTO HBA's and Operating Systems.

1 ATTO Celerity MultiPath Director™ Overview

Workstation-based applications are evolving and demanding more sophisticated, feature-rich storage solutions. These high-end professional environments require multiple users to share access to large amounts of storage while maintaining the high-throughput performance rates required for the rich media and content applications they serve. The media and entertainment industry is undergoing a transformation to all digital content and accelerated workflow requirements. Digital media professionals are driving the needs for more sophisticated configurations, which require multiple users to share a "pool" of storage for continuous access to large amounts of data.Today, multiple workstations needing access to storage rely on fragmented storage solutions that don't fulfill the high-availability, high-performance demands of rich media applications in shared SAN environments.

When organizations lose access to their storage due to hardware malfunctions, basic maintenance or changes to the storage infrastructure, the cost of storage downtime can add up quickly. Businesses concerned with this kind of downtime are moving to higher-end storage solutions that feature redundant components with multiple connection paths between the workstations and the storage.

Windows and Linux operating systems include general support for redundant storage network infrastructures, or multipathing. The Apple operating system also provides some capabilities for high-availability systems, although they tend to lack management capabilities.

ATTO Technology has developed a feature-rich driver for its Celerity Fibre Channel host bus adapters that, when paired with enterprise-class storage systems, provides a completely redundant, high-performance storage solution for workstation applications in digital video, audio and content-creation environments.

ATTO's MultiPath Director is the ONLY storage connectivity solution that allows you to manage multiple paths between your Windows®, Linux and Mac® workstations, servers and high-end storage systems.

MultiPath Director provides a high-performance collaborative workflow solution that keeps digital media content available in real time for multiple users. This allows for significant gains in productivity, minimal downtime and uninterrupted access to enterprise-class storage media. Providing multiple redundant paths to storage with load balancing and failover capabilities, ATTO's Fibre Channel host bus adapters with MulitPath Director are the premier choice for reliable and robust SAN connectivity.

Exhibit 1-1 Typical MultiPath Director Configuration.

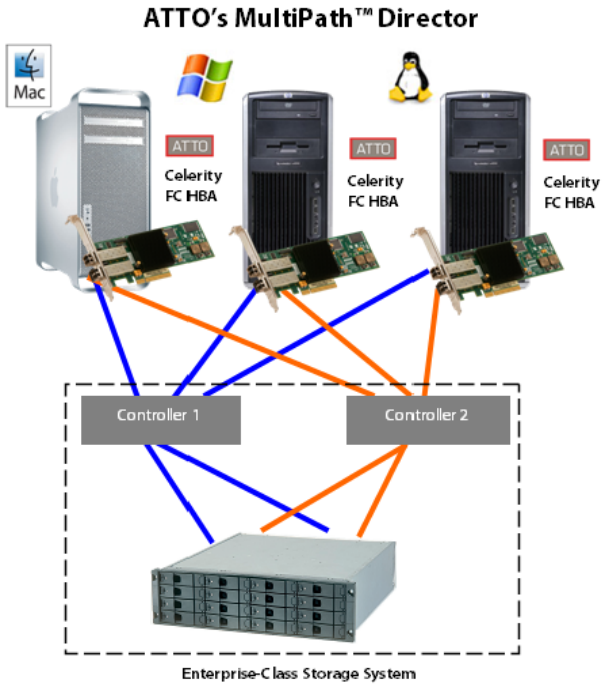

1.1 MultiPath Director Solution

- RAID protects against a drive failure
- Multiple RAID controllers protect against a controller failure
- Multiple connections to each RAID controller protects against a port, cable or SFP (optical transceiver) failure
- Multiple ATTO Celerity Fibre Channel host adapters protects against a host adapter failure
- Multiple paths to the same target allows I/O to be sent down each of these paths, effectively increasing performance

1.2 MultiPath Director Benefits

High-Availability Connectivity to Storage - Automatic path failover and failback provides uninterrupted access to storage.

Superior Performance - Load balancing increases overall system performance by using more than one Fibre Channel path to transfer data.

Improved Productivity - Workstations can share storage, enabling collaborative workflow and real-time access to content to meet project deadlines.

Flexible Connectivity - Ability to mix Windows, Linux and Mac workstations in heterogeneous environments.

Simplified Management - ATTO Configuration Tool simplifies administration, troubleshooting and management of multiple paths to storage.

1.3 Supported ATTO Celerity Fibre Channel Adapters

The table below (1.3-1) outlines the ATTO Celerity Fibre Channel Host Adapters that currently support MultiPath Director for various operating systems. Please contact your storage manufacturer for currently supported Celerity Fibre Channel Host Adapters and driver versions. All adapters may not be supported with your particular storage array.

Exhibit 1.3-1 Supported Celerity Fibre Channel Host Adapter Models

Product Features	Celerity Host Adapters					
	FC-16E4	FC-16E2	FC-16E1	FC-84EN	FC-82EN	FC-81EN
Fibre Channel Ports	4	2	1	4	2	1
Maximum Data Rate	16Gb	16Gb	16Gb	8Gb	8Gb	8Gb
Maximum Transfer Rate	6400 MB/s	3200 MB/s	1600 MB/s	3200 MB/s	1600 MB/s	800 MB/s
Bus Type	PCIe 3.0	PCIe 3.0	PCIe 3.0	PCIe 2.0	PCIe 2.0	PCIe 2.0
Bus Characteristics	x8	x8	x8	x8	x8	x8
Optical Interface	SFP+LC	SFP+LC	SFP+LC	SFP+LC	SFP+LC	SFP+LC
Maximum Cable Length	150m - 8Gb 100m - 16Gb	150m - 8Gb 100m - 16Gb	150m - 8Gb 100m - 16Gb	300m - 2Gb 150m - 4Gb 50m - 8Gb	300m - 2Gb 150m - 4Gb 50m - 8Gb	300m - 2Gb 150m - 4Gb 50m - 8Gb
Low-Profile Form Factor	✓	✓	✓		✓	✓
Advanced Data Streaming (ADS™)	✓	✓	✓	✓	✓	✓
Software RAID Support	✓	✓	✓	✓	✓	✓
Developer's Kit (Target Mode & API)	✓	✓	✓	✓	✓	✓
Windows®	✓	✓	✓	✓	✓	✓
Linux® (Red Hat, SUSE)	✓	✓	✓	✓	✓	✓
Mac® OS X	✓	✓	✓	✓	✓	✓
VMware® ESX Server	Fall 2012	Fall 2012	Fall 2012	✓	✓	✓
RoHS Compliant	✓	✓	✓	✓	✓	✓

2 Installing the Celerity HBA Driver

For hardware installation, please see the ATTO Celerity Fibre Channel Host Adapter manual.

Note

*The MultiPath Director driver is **NOT** included on the CD that ships with standard ATTO products.*

2.1 Driver Installation - Mac OS X (Supports 4, 8 and 16Gb)

1. Download the appropriate MultiPath Director driver file from www.attotech.com/downloads.html. Make sure it is the version qualified with the storage manufacturer you are connecting to.
2. Double-Click the **osx_drv_celerity_xxxMP.dmg** (4 Gig Fibre Channel), **osx_drv_celerity8_xxxMP.dmg** (8 Gig Fibre Channel) or **osx_drv_celerity16_xxxMP.dmg** (16 Gig Fibre Channel)file (xxx is the version number of the driver). This will create and launch the disk image.
3. Double-click **ATTOCelerityFC_xxxMP.pkg** or **ATTOCelerityFC8_xxxMP.pkg** to begin installation.
4. Click **continue** to agree to the “Software License Agreement”.
5. Select the “Destination drive (boot drive)” and select **continue**.
6. Select **Install**.
7. The driver will now automatically install. When complete the following screen will appear, click **restart** to reboot you machine.

2.2 Driver Installation - Windows (Supports 8 and 16Gb)

Driver installation for Windows operating systems is a two step process

2.2.1 Installing Celerity Drivers

1. Download the appropriate storport driver file from www.attotech.com/downloads.html. Make sure it is the version qualified with the storage manufacturer you are connecting to.
2. Double-click the **winyyy_drv_celerity8_xxx.exe** file for 8Gb or **winyyy_drv_celerity16_xxx.exe** file for 16Gb (xxx is the version number of the driver and yyy is the processor architecture, eg. x86 or x64) to extract the driver files.
3. When the driver files are extracted, the installer will be launched automatically. Click **Install** to install the driver.

2.2.2 Installing MultiPath Director Drivers

1. Download the appropriate MultiPath Director driver file from www.attotech.com/downloads.html. Make sure it is the version qualified with the storage manufacturer you are connecting to.
2. Double-click the **winyyy_drv_mpdirector_xxx.exe** file (xxx is the version number of the driver and yyy is the processor architecture, eg. x86 or x64) to extract the driver files.
3. When the driver files are extracted, the installer will be launched automatically. Follow the on-screen instructions until the installation completes.

2.3 Driver Installation - Linux (Supports 8 and 16Gb)

1. Download the appropriate MultiPath Director driver file from *www.attotech.com/downloads.html*. Make sure it is the version qualified with the storage manufacturer you are connecting to.
2. Copy the **filename.tgz** file (xxx is the version number of the driver) to a temporary folder.
3. Open a terminal window and change to the location of the tgz file from step 2.
4. Extract the file using the command **tar xfz <filename.tgz>**
5. Change to the directory created above, then run the installer script **./install.sh**

3 Installing the ATTO Configuration Tool

Download the ATTO Configuration Tool from www.attotech.com/downloads.html.

Note

The default path of the ATTO Configuration Tool is /Applications/ATTO Configuration Tool.

3.1 Configuration Tool Installation - Mac OS X

1. Double-click **osx_app_configtool_xxx.dmg** file (xxx is the version number of the application).This will create and launch the disk image.
2. Double-click **ConfigTool_xxx** file (xxx is the version number of the application).
3. Click next, to agree to the “Software License Agreement”.
4. Select the “Default Folder” and click next.
5. Choose “Full Installation” and click next.
6. The ATTO Configuration Tool will now install.
7. The following screen will appear when the install has completed.

3.2 Configuration Tool Installation - Windows

1. Download the ATTO Configuration Tool from *www.attotech.com/downloads.html*
2. Double-click the **win_app_configtool_xxx.exe** file (xxx is the version number of the application).
3. Follow the on-screen instructions until the installation completes.

3.3 Configuration Tool Installation - Linux

1. Verify in a terminal window that the installed Java Virtual Machine (Java) is from Sun by typing java -version. The GNU version JVM which may be installed by default does not work with the installer. See section 3.3.1 Configuring Java for Configuration Tool for help.
2. Download the ATTO Configuration Tool from *www.attotech.com/downloads.html*
3. Double-click the **lnx_app_configtool_xxx.bin** file (xxx is the version number of the application).
4. Follow the on-screen instructions until the installation completes.

3.3.1 Configuring Java for Configuration Tool

1. Download and install the latest version of Java from the SUN website (http://www.java.com). ATTO recommends installing the BIN version over the RPM package as the BIN procedure is easier to use. The following instructions are recommended by ATTO for installing the JAVA BIN package:
   a. Make java directory
      i. mkdir -p /usr/java
   b. Save java BIN file to this directory.
   c. Change to java directory in terminal
      i. cd /usr/java
   d. Determine if BIN file is executable.
      i. ls -l
      ii. Look for “x” in the file characteristics.
   e. If BIN file not executable, make it so…
      i. chmod a+x <BIN_FILE_NAME>
   f. Execute BIN file and follow prompts to install Java.
      i. ./<BIN_FILE_NAME>

Note

Refer to SUN installation instructions for more details or troubleshooting.

2 When Java installation is complete, type the following in terminal:

alternatives --install <LINK> <NAME> <PATH> <PRIORITY>

where:
<LINK> = /usr/bin/java
<NAME> = java
<PATH> = Path to Java installation (i.e. /usr/java/jre1.6.0-14/bin/java)
<PRIORITY> = 0

If this installs successfully, it will just return to a command prompt with no other messages.

Note

In SUSE 11 "alternatives" is now "update-alternatives". All other ATTO supported Linux versions use "alternatives".

3 Type the following in terminal:

alternatives --config java

You will be prompted to enter the number of the java installation you would like to use. Find the version you just installed and enter the number associated with it.

For example:

```
[root]# alternatives --config java
  There are 2 programs which provide 'java'.

   Selection    Command
-----------------------------------------------
*+ 1           /usr/lib/jvm/jre-1.4.2-gcj/bin/java
   2           /usr/lib/jvm/jre1.6.0/bin/java

Enter to keep the current selection[+], or type selection number: 2
[root]#
```

4 Download and Install the Configuration Tool (go back to Step 2 in secton 3.3 Configuration Tool Installation - Linux)

4 Sample MultiPath Director Configuration

The figure below (4-1 Cross Zoned Multipathing Configuration) shows an example of a cross-zoned multipathing configuration.

Exhibit 4-1 Cross Zoned Multipathing Configuration

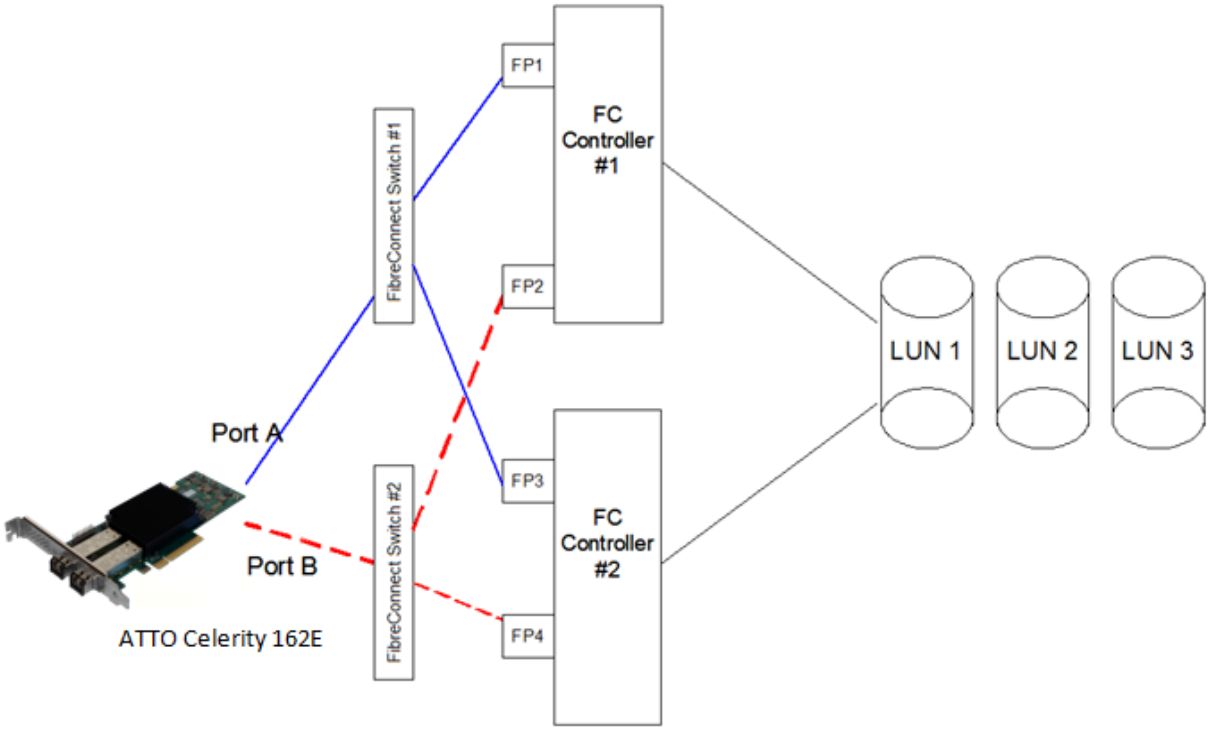

There are (2) paths from the host, one to each of the switches. Each switch has (2) paths to the target, one to each of the RAID controllers.

This configuration provides (4) paths into the storage from two independent switches and adapter ports. This is one of the most common methods of multipathing in the industry.

- Celerity 82EN/162E Port A ---- Target Port FP1 and Target Port FP3 via Switch #1
- Celerity 82EN/162E Port B ---- Target Port FP2 and Target Port FP4 via Switch #2

If a connection between the host adapter port and either of the switches were to fail (Blue/Port A), all IO would be transferred to the opposite host adapter/switch (Red/Port B) path. Likewise, if a connection on the back-end (switch to storage) were to fail, IO would be redirected automatically to the available paths. This provides the user simple failover and load balancing.

There are several different ways to configure:

- Automatic (via ALUA / most storage arrays)
- Manual by Target (most simple manual configuration)
- Manual by LUN (greatest level of manual configuration)

Note

Automatic is the recommended mode for reliability and performance.

5 Viewing Multipathing Information - Paths Tab (Status)

*One of the strengths of the ATTO MultiPath Director driver is the ability to monitor and manage the multiple connections between the host and the storage. The ATTO Configuration Tool will show all associated paths to each target or LUN as well as the status of that path. It can also be used to track statistics for data transferred across each path. The **Paths** tab in the multipathing window displays connection path information from the host adapter ports to each individual target port or LUN in the storage system.*

5.1 Paths (Target Basis)

Click on the Target in the device tree within the window on the left to display the Target “Paths” multipathing window on the right. In most cases, the Target “Paths” information relates to the connection between the host adapter and the RAID Controller Ports (Target). Most RAID Controllers have more than one port - most storage systems use at least two RAID Controllers. The result is multiple paths to a Target for a typical configuration. The “Paths” window in *5.1-1 Monitoring Paths to a LUN* below shows four connection paths to the Target labeled “Storage System”. This represents a system with two RAID controllers, each with two input connections (see Exhibit4-1 on page 7). For this storage system, each connection port on the RAID Controllers is identified by a unique World Wide Name.The “Paths” window identifies how each path is configured for read and write operations and the current state of that path.

Exhibit 5.1-1 Monitoring Paths to a Target

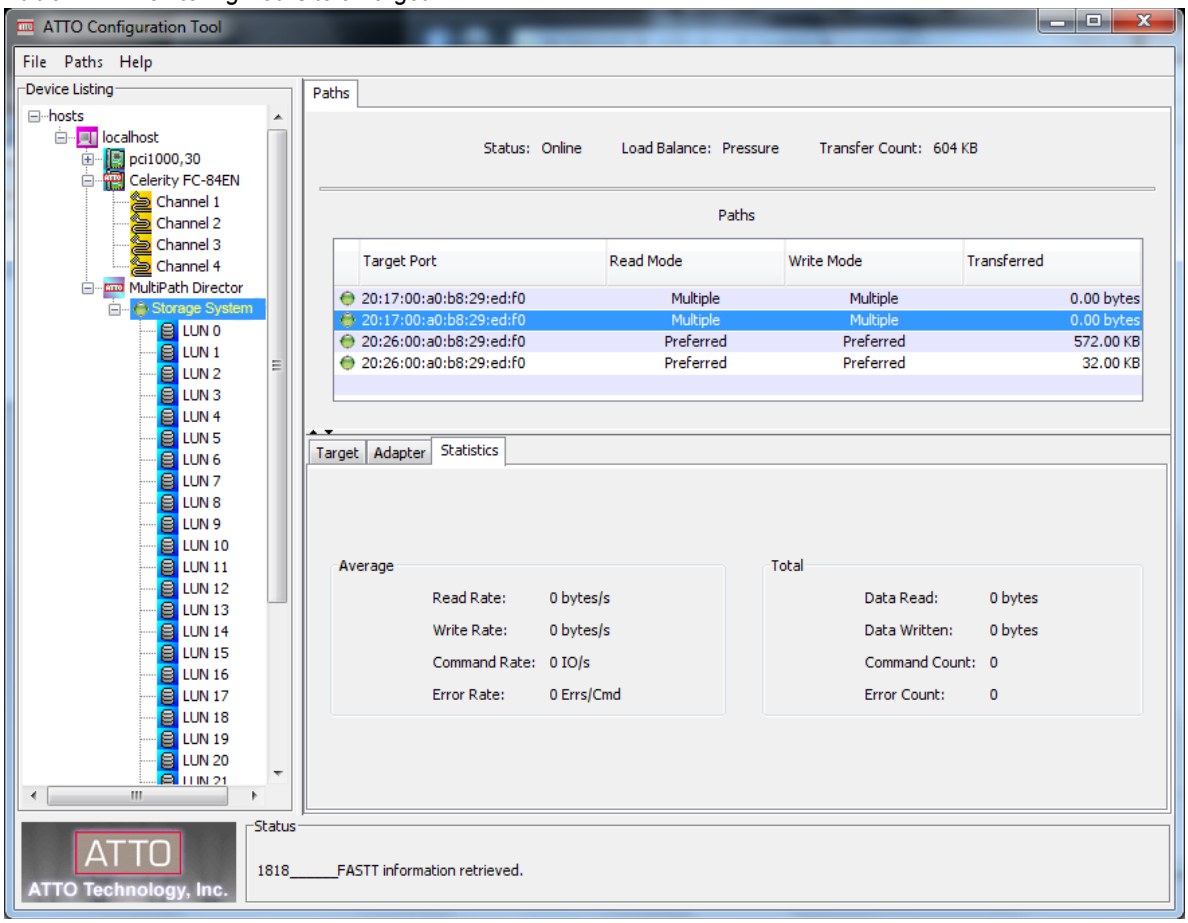

Overall Status at a Target Level

The icon next to the device in the left hand pane indicates the overall multipathing status for the target (NOT each LUN).

- A single path is connected to your device
- Multiple paths are connected to your device
- Multiple paths were connected to your device, but one or more have failed
  NOTE: This will only be visible if a configuration was previously saved (See section 5.4)
- Multiple paths were connected to your device, but all have failed
  NOTE: This will only be visible if a configuration was previously save. (See section 5.4)

Path Status at Target Level

The icon next to the Target Port in the right hand pane under the "Paths" tab when a target is selected indicates the connection status for the target port.

- Paths are connected to the target and are functioning normally
- Paths were connected to the target and are no longer connected
  NOTE: This will only be visible if a configuration was previously saved (See Section 5.4)

Additional Management and Monitoring Capabilities at the Target Level

The following additional functions are available within the "Paths" window at a Target level. (Refer to *Exhibit 5.3-1 on page 13*).

Status

This represents the status of the sum of all paths connected to the highlighted Target device.

Online – All paths connected to a device are connected and operating properly.

Degraded – One or more paths have failed (if the configuration has been saved).

Changed – One or more paths have been added since the configuration was last saved.

Offline – All paths to a particular Target have failed (if the configuration has been saved) and connectivity has been lost.

Load Balance

The load balancing policy for a Target. See Chapter 6 page 17 for more info.

Transfer Count

The total number of bytes transferred to a target. (You can reset these statistics by going to "Paths" in the Tool Bar and selecting "Reset Statistics"). These stats can be useful in determining if you are properly utilizing the paths.

Target Port

The WWPN (World Wide Port Name), or unique Fibre Channel Address for each Target port on the RAID controllers.

Read Mode

Indicates the selected path's read policy.

Preferred – This will be the first path that the driver will attempt to read IO through.

Alternate – This path will be used for read operations if the preferred path has failed.

Multiple – All LUN paths behind the target are not set the same. Path A (Preferred, Alternate, Disabled) for LUN 0 can be preferred and Path A for LUN 1 could be an alternate; this would result in multiple read mode.

Disabled – The associated path will not be used for read operations.

Write Mode
Indicates the selected path's write policy.

Preferred – This will be the first path that the driver will attempt to write IO through.

Alternate – This path will be used for write operations if the Preferred path has failed.

Multiple – All LUN paths behind the target are not set the same. Path A (Preferred, Alternate, Disabled) for LUN 0 can be preferred and Path A for LUN 1 could be an alternate; this would result in multiple write mode.

Disabled – The associated path will not be used for write operations.

Transferred
The total number of bytes transferred on the associated path.

5.2 Paths (LUN Basis)

The LUN "Paths" display helps to manage and monitor the connections to one or more virtual LUNs created by the RAID Controller. Click on a LUN in the device tree within the window on the left to display the expanded window on the right. Note that the Configuration Tool provides three options when viewing LUNs: Basic Info, Flash and Paths. Click on the "Paths" option in the right half window to display the path details for each LUN. With most RAID Controllers, the same LUN is usually assigned to more than one controller, resulting in multiple paths to each LUN. The display will list out each path, along with its current status for the selected LUN.

Exhibit 5.2-1 on page 11 shows an example of LUN Path information. In this case, there are 4 paths to LUN 1. The WWN for each Target Port on the RAID controllers is listed to show the physical paths to each LUN. Each path can have the following status:

Path Status at LUN Level

The icon next to the Target Port in the right hand pane under Paths when a LUN is selected indicates the connection status for the LUN.

- Path connected to selected LUN is functioning normally and available for I/O.
- Path connected to the selected LUN has failed.
  NOTE: This will only be visible if a configuration was previously saved. (See section 5.4)
- Path to the selected LUN is set to Alternate for Read and Write functionality.
- Path to the selected LUN is Logically DISABLED for Read and Write functionality.

Exhibit 5.2-1 Monitoring Paths to a LUN

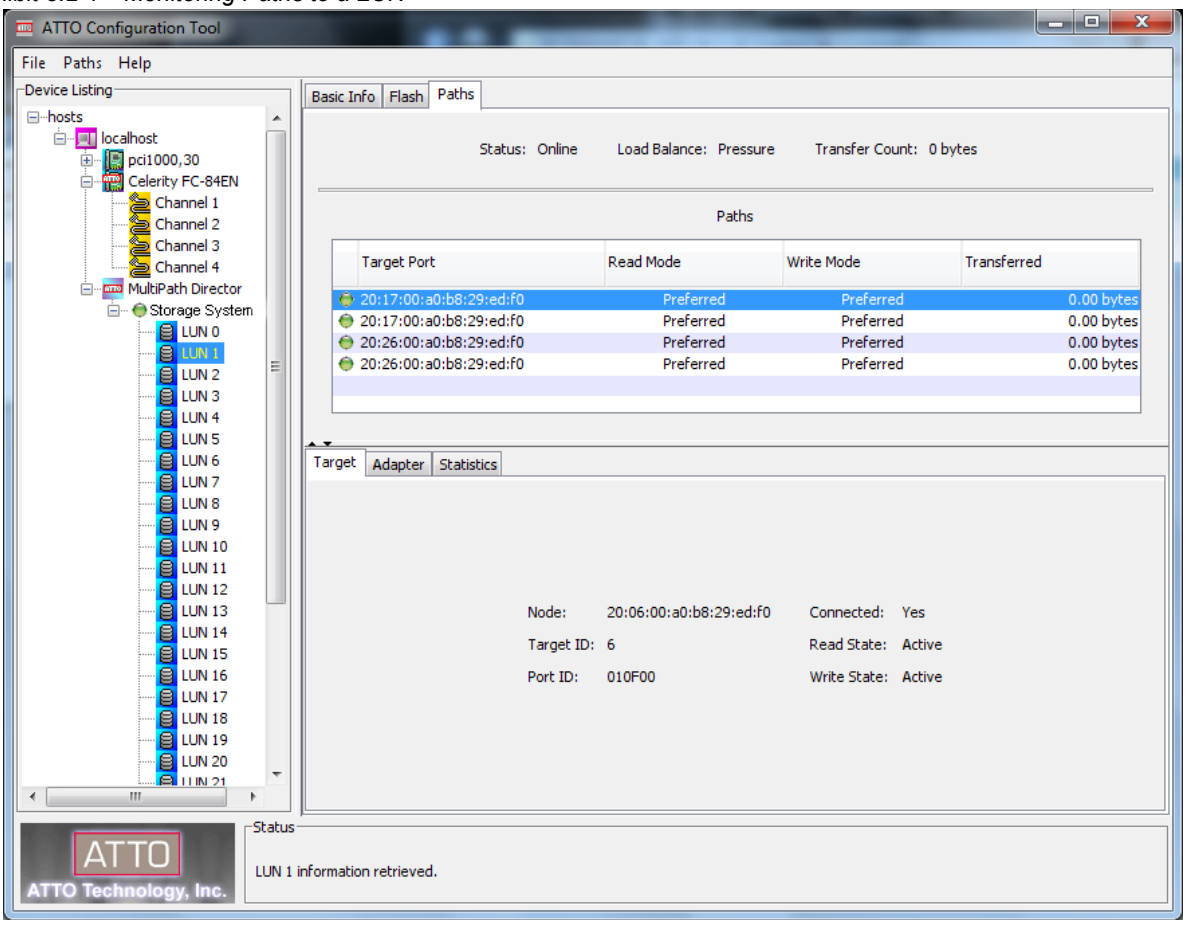

Additional Management and Monitoring Capabilities at the LUN Level

The following additional functions are available within the "Paths" window at a LUN level. Refer to *5.2-1 Monitoring Paths to a LUN*.

Status

This represents the status of the paths connected to the highlighted LUN.

Online - All paths connected to a device are connected and operating properly.

Degraded - One or more paths have failed (if the configuration has been saved).

Changed - One or more paths have been added since the configuration was last saved.

Offline - All paths to a particular LUN have failed (if the configuration has been saved) and connectivity has been lost.

Load Balance

The load balancing policy for a LUN. See Chapter 8 page 17 for more info.

Transfer Count

The total number of bytes transferred to a LUN. (You can reset these statistics by going to "Paths" in the Tool Bar and selecting "Reset Statistics"). These stats can be useful in determining if you are properly utilizing the paths.

Target Port

The WWPN (World Wide Port Name) or unique Fibre Channel Address for each Target Port on the RAID controllers.

Read Mode

Indicates the selected path's read policy.

Preferred - This will be the first path that the driver will attempt to read IO through.

Alternate - This path will be used for read operations if the Preferred path has failed.

Disabled - The associated path will not be used for read operations.

Write Mode
Indicates the selected path's write policy. Eg. Preferred, Alternate, Multiple (for targets ONLY) or Disabled.

Preferred - This will be the first path that the driver will attempt to write IO through.

Alternate - This path will be used for write operations if the Preferred path has failed.

Disabled - The associated path will not be used for write operations.

Transferred
The total number of bytes transferred on a particular path.

5.3 Path Details

Select a Target or LUN in the left hand pane of the Configuration Tool and click on the “Paths” tab in the top of the right hand pane. This will display additional information for the path on a Target or LUN basis in the lower half of the window. The following information is available in three tabs:

1. Target

When a target or LUN is selected, this window will display information for the storage device's target port that is connected to the path. (Refer to Exhibit 5.3-1 on page 13)

Node
World Wide Node Name of the Target Port. This is useful in determining which target Node you are connected to.

Target ID
Internal ID used by the ATTO driver

Port ID
Fibre Channel Port ID

Connected
Link Status (Yes, No). Used to identify if the device is logged into the ATTO adapter.

Read State
Status of Read Policy (Active, Alternate, Disabled, Failed, Multiple)

Active - The associated path to the target will be used for read operations.

Alternate - The associated path to the target will only be used as a read alternate in the event there are no active paths.

Disabled - The associated path will not be used for read operations.

Multiple - Special case only available if the target is selected in the left hand pane. The associated paths have several states (active, alternate or disabled) for that target. For example Target 0 could have Path A for LUN 0 as active and Path A for LUN1 as alternate; this would result in a multiple state for that target.

Failed - If the associated path was saved, this path has failed for reads.

Write State
Status of Write Policy (Active, Alternate, Disabled, Failed, Multiple)

Active - The associated path to the target will be used for write operations.

Alternate - The associated path to the target will only be used as a write alternate in the event there are no active paths.

Disabled - The associated path will not be used for write operations.

Multiple - Special case only available if the target is selected in the left hand pane. The associated paths have several states (active, alternate or disabled) for that target. For example Target 0 could have Path A for LUN 0 as active and Path A for LUN1 as alternate; this would result in a multiple state for that target.

Failed - If the associated path was saved, this path has failed for writes.

Exhibit 5.3-1 Target Information

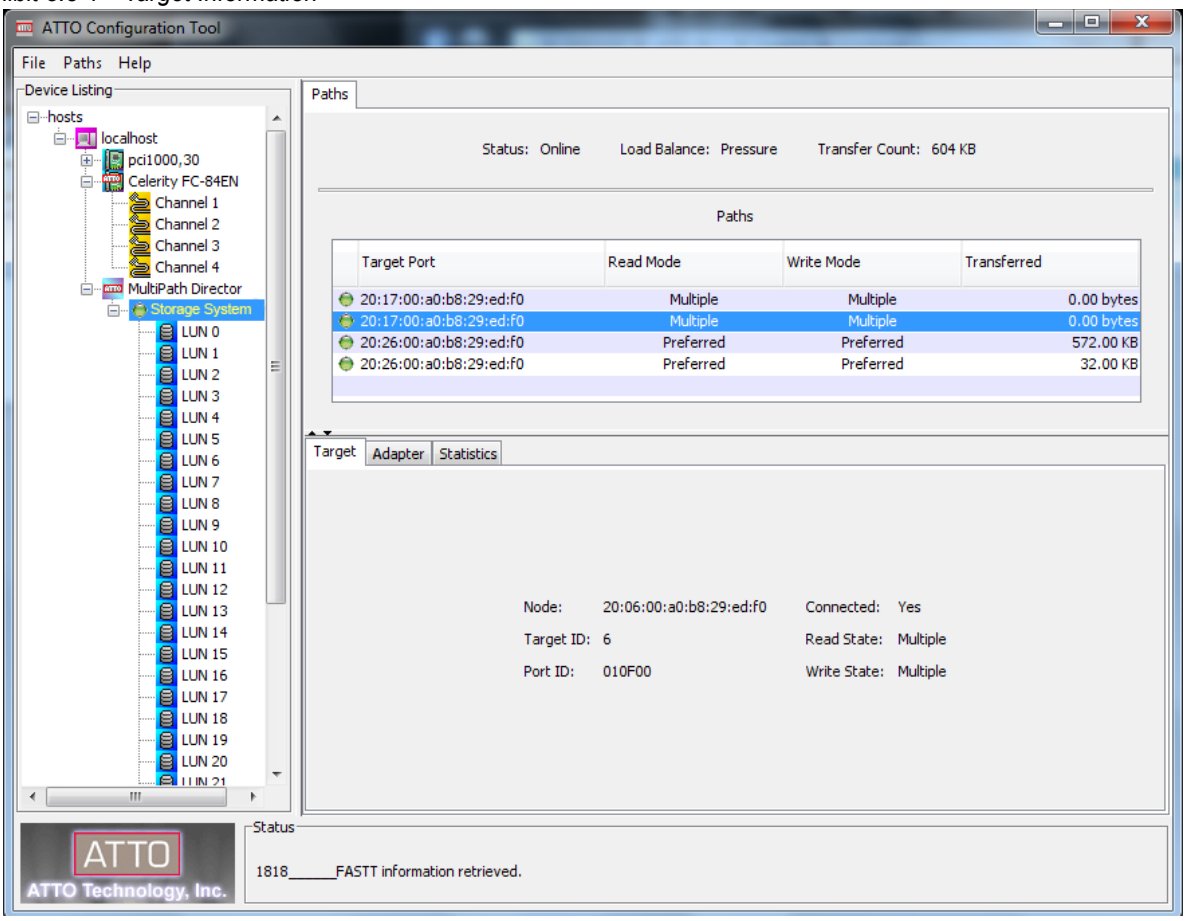

2. Adapter

When a path is selected it displays information for the Celerity host adapter channel that is connected to that path. (Refer to Exhibit 5.3-2 on page 14)

Node - World Wide Node Name of the Celerity FC Adapter port.

Port - World Wide Port Name of the Celerity FC Adapter port.

Port ID - Fibre Channel Port ID of the Celerity FC Adapter port.

Link - Status of Fibre Channel Link (Up, Down).

PCI Location - Location of the Fibre Channel Adapter in the PCI bus of the computer.

Domain ID - Number of the Fibre Channel Domain.

Topology - Fibre Channel Topology (Point to Point, Arbitrated Loop, Unknown).

Link Speed - Fibre Channel Data Rate (16Gb/s, 8Gb/s, 4Gb/s, 2Gb/s,1Gb/s).

Exhibit 5.3-2 Monitoring Paths to a LUN

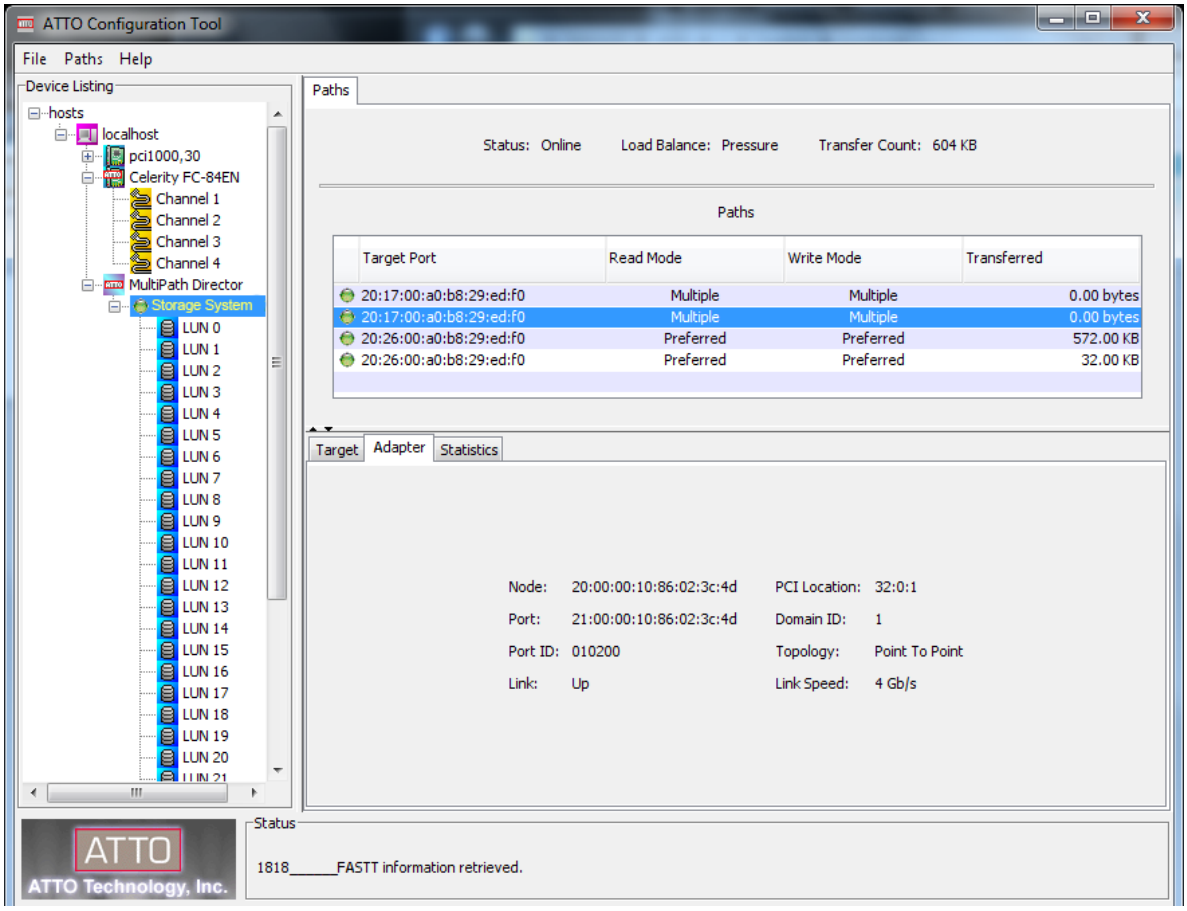

3. Statistics

When a Target or LUN path is selected, this window displays information regarding data transferred down that path. (Refer to Exhibit 5.3-3 on page 15)

Average

Read Rate – Number of bytes read per second for a particular path.

Write Rate – Number of bytes written per second for a particular path.

Command Rate – Number of commands sent down a particular path in IO's per second.

Error Rate – Error rate in errors per command for a particular path.

Total

Data Read – Total amount of data in bytes read for a particular path.

Data Written – Total amount of data in bytes written for a particular path.

Command Count – Total number of commands sent down a particular path.

Error Count – Total number of errors for a particular path.

Exhibit 5.3-3 Monitoring Paths to a LUN

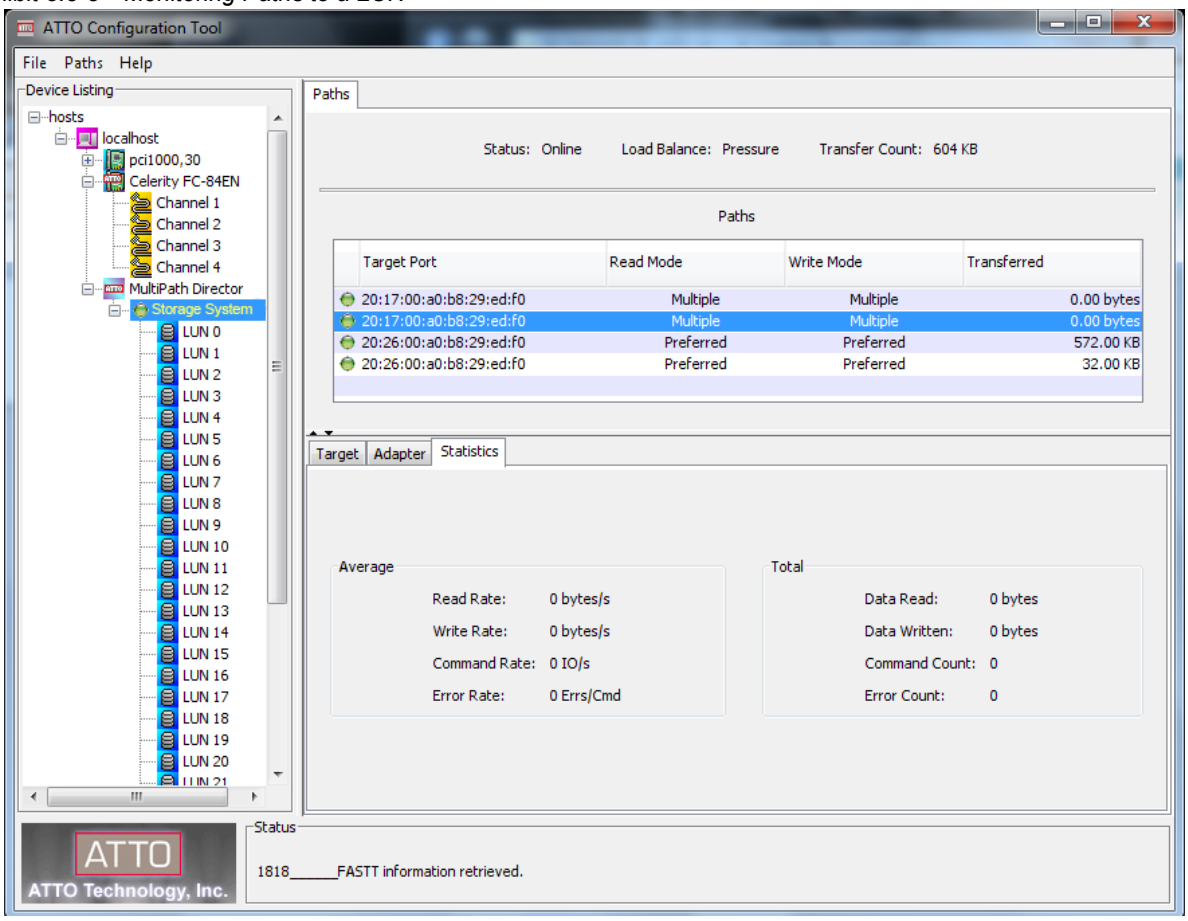

5.4 Saving your Configuration

Note

In order to see the associated states of each path (degraded, failed, etc), it is necessary to save your configuration after initial setup. It is also recommended you save your setup after any changes to your SAN.

To save your configuration:

1 Highlight the Target in the left pane.

2 Select “Paths” from the toolbar.

3 Select “Save Configuration” or CTRL + S (or COMMAND + S in Mac OS).

4 You will be prompted to save current configuration - select “Yes”.

Your multipathing configuration is now saved.

Note

If a LUN is selected in the left pane, only the paths for that particular LUN will be saved.

5.5 Paths Menu

The following functionality can be accessed via the "Paths" menu in the Toolbar, or by right clicking (CTRL + clicking in Mac OS) in the Paths window. Some items also have keyboard shortcuts.

- The "Setup" option in the Paths menu will allow you to configure your paths or load balancing options. The "Setup" option may not appear with some storage systems. In these cases, the optimal path configuration has been predetermined and will automatically be configured (see *Chapter 8 page 19*). Users can still select the appropriate "Load Balancing" policy when the paths are automatically configured.
- The "Save Configuration" option will save all paths at any given point in time when the Target is selected. This will allow the user to determine if there has been a link failure or other errors.

Note

It is recommended that you save your configuration after you verify your multipathing setup is operating properly.

- The "Delete Configuration" will delete the saved configuration. This can be useful if devices have been removed from your SAN or changed in your SAN.
- The "Locate" button will blink the particular FC adapter's port LED (very useful troubleshooting tool). You can also locate via CTRL + L (or COMMAND + L in Mac OS).
- Statistics can be reset to '0' by choosing "Reset Statistics".
- Paths can be manually refreshed by choosing "Refresh" or CTRL + R (or COMMAND + R in Mac OS). The Configuration Tool will periodically (every 10 seconds) refresh the path statistics for all devices. If the connection status for a path changes, the Configuration Tool will automatically refresh the display.

Exhibit 5.5-1 Paths Menu

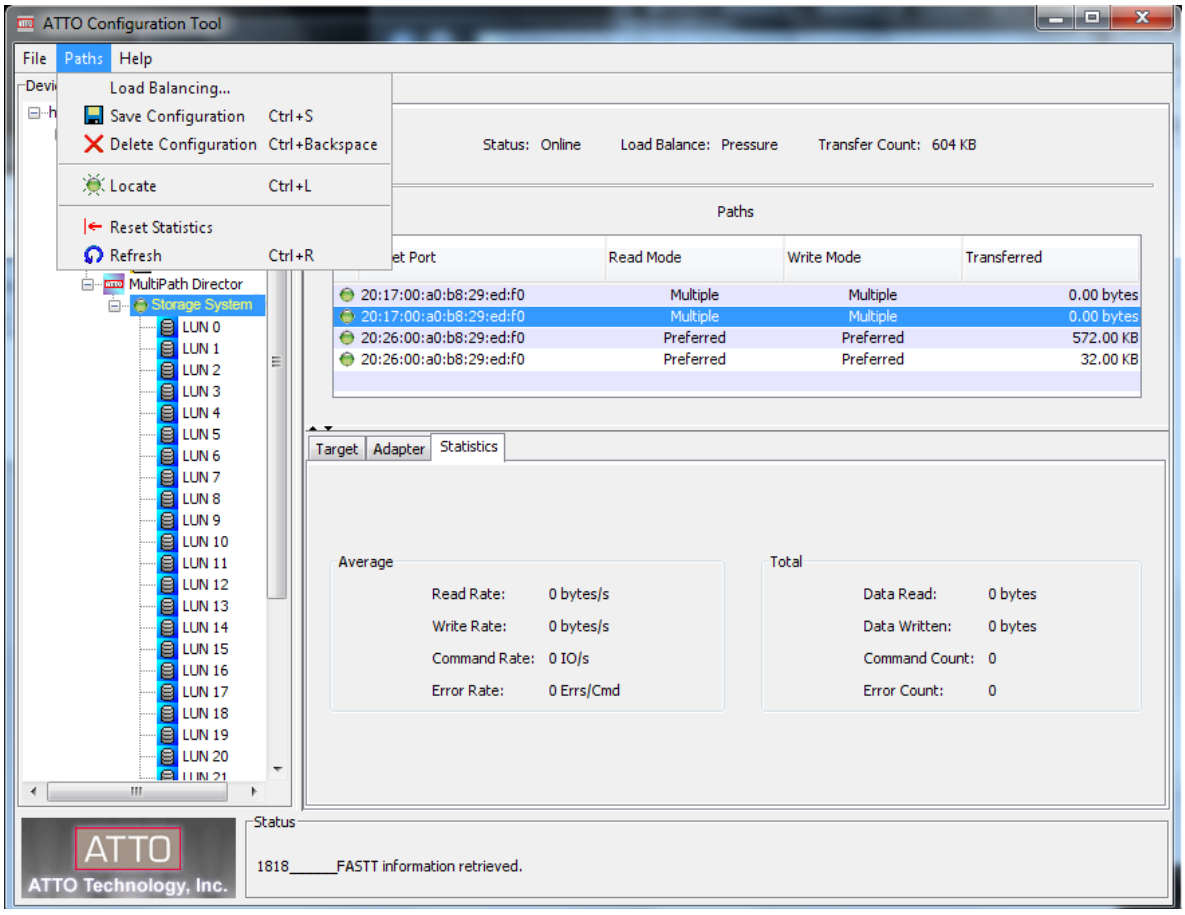

6 Load Balancing Policies and Configuration

In addition to failover, the MultiPath Director driver also supports load balancing. This is where the host adapter driver can utilize more than one Fibre Channel link to move data to/from a target or a LUN. The result is a considerable increase in performance. The ATTO Celerity Host Adapter will load balance whenever there are two or more paths to a LUN and at least two of these paths are "preferred". Load balancing policies (algorithms) can be set on a per Target or a per LUN basis, giving the end user more control of how performance is balanced across RAID controllers and workstations in a SAN environment. ATTO currently supports three load balancing policies (Pressure, Round Robin, and Queue Depth).

Pressure (default)
The path with the least number of bytes currently being transferred is selected for I/O. Pressure based load balancing is domain based. Path selection is based on the domain (adapter channel) with the fewest bytes being transferred. For example, an application sends a 1 MB I/O to LUN 1 that goes down path 1. If the application then issues two 512KB I/Os to LUN 2, they both will be transferred on path 2. If a target based approach was taken, 1.5 MB of data would have been transferred on one path while 512KB would have went on the other – not a good use of available bandwidth. This load balancing algorithm is designed for performance when large blocks are transferred. It usually results in the best performance for most applications.

Queue Depth
The path with the fewest I/Os outstanding will be selected for the next I/O. This policy is strictly target based.Consider an application that simultaneously issues 10 various size I/O's to the target. Assuming there are two paths to the target, 5 I/O's will go down one path and 5 will go down the other. The size of each I/O does not matter in this case. This load balancing policy is designed for those storage devices that have resource limitations on each target port. If a dual ported device can only handle 10 I/Os on each port, an application can have a queue depth of 20 and be guaranteed the target can process all requests immediately regardless of I/O size.

Round Robin
The least used path is selected for I/O. This policy is also target based. Like Queue Depth, an I/O count is used for path selection; however, the count is cumulative for all I/O in the current configuration (NOT the number of currently queued I/O's). For example, if an I/O operation starts and finishes on one path to a dual ported device, the second path is used when the next I/O operation starts. The round robin policy exercises all paths when the application queue depth to a target is 1. In other words, it ensures that all hardware devices are used evenly. The other policies select the same path repeatedly when the queue depth is 1. The primary reason for using this policy is to exercise all paths when the application queue depth to a target is one. The other policies will always select the same path when the queue depth is one.

6.1 Adjusting Load Balancing Policies

Load Balancing can be configured on a target or LUN basis. If you select the target in the left hand pane, you will set the load balancing policy for all LUNs associated with that target. If you select a LUN you will only set the load balancing policy for that particular LUN. All paths associated with a LUN or target must share the same load balancing policy.

1 Select the "Target" or "LUN" in the left hand pane of the Configuration Tool.
2 Select "Paths" in the toolbar. You can also right click (or CTRL - click in Mac OS) inside the paths window.
3 Select "Setup" or "Load Balancing".
4 Select a policy. (Pressure, Round Robin, or Queue Depth)
5 Select "Finish".


Note
It is not necessary to save the configuration to save load balancing settings.

7 ALUA – Asymmetric Logical Unit Access

ALUA is the behavior defined in the SCSI specification that describes how storage devices with multiple ports are to behave. Asymmetrical Logical Unit Access occurs when the access characteristics of one port differ from another. All ports that have the same behavior are referred to as target port groups. Typically, they are referred to as controllers on storage devices. Among other things, each target port group has an asymmetric access state. The MultiPath Director driver will use this state to automatically configure path modes for optimal performance.The ALUA state can be one of the following (with the corresponding path mode in parentheses):

Active/optimized (Preferred)
The path can execute all commands supported by the device with optimal performance.

Active/non-optimized (Alternate)
The path can execute all commands supported by the device with less than optimal performance.

Standby (Alternate)
The path can only execute a subset of commands supported by the device. This typically does not include I/O commands (ex. Reads and Writes). The path can be activated to allow I/O if needed, but usually another active path is reverted to standby at the same time.

Unavailable (Disabled)
The path can only execute a subset of commands supported by the device and it is not available for I/O (ex. Reads and Writes). This state is typically used when a hardware failure has occurred or the device has been placed in a maintenance or service mode.

ALUA allows for target port groups to transition from one state to another. These transitions can be determined by the storage device (implicit transitions) or by the host adapter (explicit transitions). The ATTO MultiPath Director driver supports both of these models.

Note

The preferred mode is reported to the driver from the target array.

Implicit transitions are automatically detected by the driver and the path modes are updated accordingly. The driver will initiate an explicit transition from the standby state to an active state only when no viable active paths to the device exist. During this process, I/O is suspended and will resume as soon as the transition is complete.

Alternate paths are activated in a tiered approach. Active/non-optimized alternates are immediately activated to take the place of a preferred path. Standby alternates are activated only when there are no remaining active paths.

Note

Saving the configuration does not override ALUA functionality.

Events that result in failover generally fall into two classes:

1 Link/protocol events:
- SFP failure
- Cable pull
- Controller power down
- RSCNs (Registered State Change Notification)
- LIPs (Loop Initialization Primitive)
- Other events that result in the target device becoming (temporarily) unavailable for I/O

2 When a target is present (logged in), commands are monitored for errors to determine when the path should be logically deactivated. For this reason, all failover is managed at a LUN level.

When a path fails:
- Active/non-optimized alternates are immediately activated to take the place of a preferred path when no preferred paths exist.
- Standby alternates are activated only when there are no remaining active paths.
- Disabled paths will never be activated.

8 Manual Configuration via the Setup Wizard

Every storage device behaves differently in terms of how it handles multiple path configurations to its LUNs. Arrays that support ALUA will most often automatically negotiate path configuration with the ATTO host adapter.

There are a few storage arrays that do not support ALUA. Users will need to manually configure each path to each LUN for these configurations.

An exception to the dynamic path configuration based on ALUA state is when a user needs to manually configure a path mode with the Configuration Tool. In this case, the path mode is never changed until the user deletes the original configuration. This is done via the “Setup” Wizard.

CAUTION

Manual Configuration will override the ALUA configuration and is NOT recommended! This is only recommended if your storage does not support ALUA. Please contact your storage manufacturer to determine if an array supports ALUA.

The following settings can be configured on a per LUN or per Target basis (depending on which is highlighted in the left hand pane). If the Target is highlighted, then **ALL** LUN’s controlled by it will be configured the same. Users can also customize settings on a per LUN basis.

- Write Mode per Target or LUN (Preferred, Alternate, or Disabled). This allows you to steer writes down a particular path.
- Read Mode per Target or LUN (Preferred, Alternate, or Disabled). This allows you to steer reads down a particular path.
- Load Balancing Policy (see *Chapter 6*) per Target or LUN (Pressure, Round Robin, or Queue Depth).

To configure the path:

1. Highlight either the Target or LUN to be configured.
2. Navigate to the Paths tab in the toolbar.
3. Choose “Setup”.
4. Notice the configurable options in the lower right pane.
5. Select the Load Balancing Policy of choice Target or LUN highlighted.
6. Choose “Next”.
7. Choose Read policy for the path highlighted
8. Choose Write Policy for the path highlighted.
9. Select “Next”.
10. Choose Load Balancing, Read, and Write Policy for each of the remaining paths.
11. When done, select “Finish”. (Clicking “Finish” automatically saves the configuration and overrides any ALUA or previously saved configuration).

9 Support Information and NVRAM Settings

9.1 Verifying and Collecting Celerity Host Adapter Info

The ATTO Configuration Tool will display "Basic" information on the ATTO Celerity Host Adapter(s) installed in the host when it is selected in the device tree within the window on the left. Once the appropriate host adapter is selected, click on the "Basic Info" tab at the top of the window on the right (See 9.1-1). The following host adapter information will be displayed:

- HBA model number
- Device Type
- Bus Type
- PCI slot number the adapter is located in
- Location and name of the Celerity driver installed
- HBA driver version

Exhibit 9.1-1 Host Adapter Basic Info

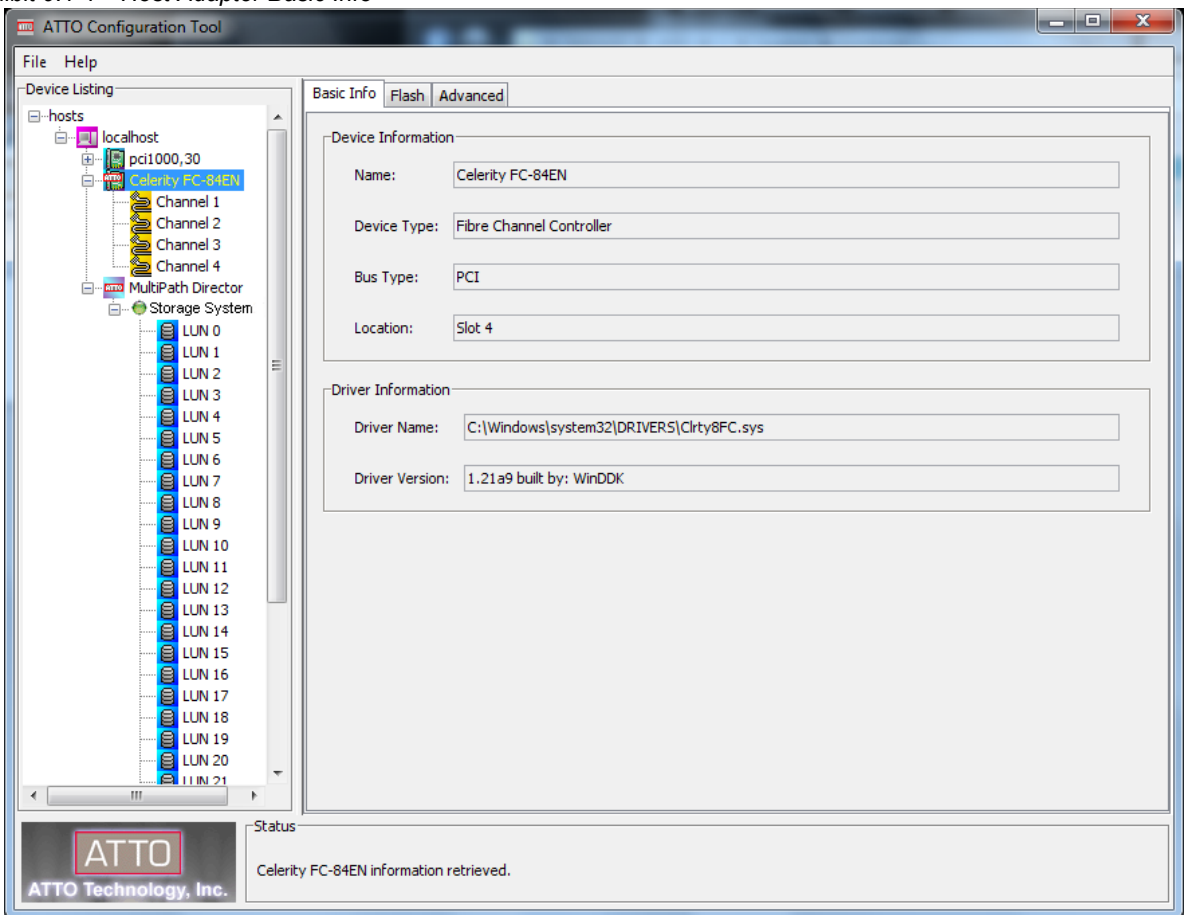

9.2 Verifying and Updating the Flash for the Celerity Host Adapter

The ATTO Configuration Tool will display the current version of the Flash installed on the ATTO host adapter(s) as well as provide a method for updating the Flash. Select the appropriate host adapter from the device tree within the window on the left. Click on the "Flash" tab at the top of the window on the right (See 9.2-1). The following flash information will be displayed:

- HBA Flash version
- Flash update information
- Flash update progress bar

Exhibit 9.2-1 Host Adapter Flash Info

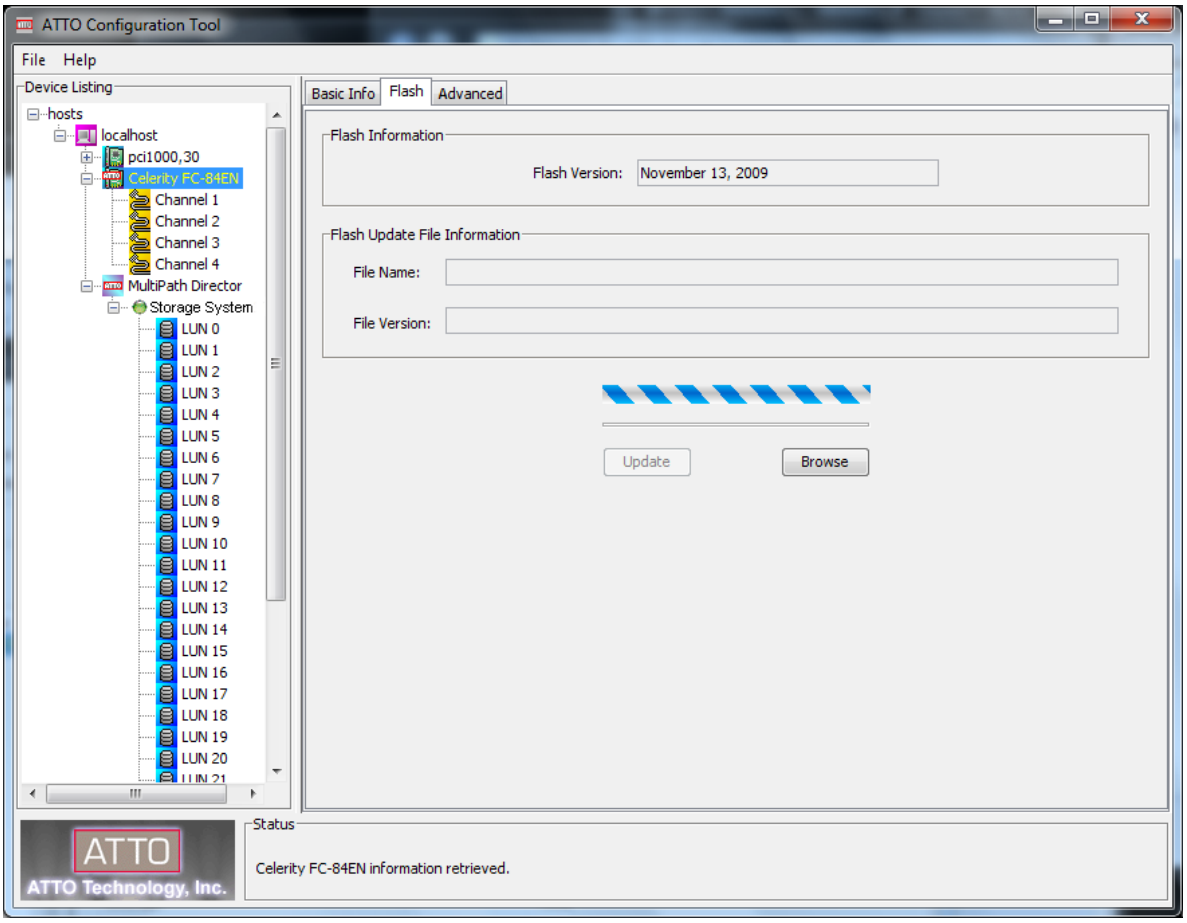

To update the Flash version of the HBA from this pane.

1 Select "Browse"
2 Locate the proper flash bundle available at the ATTO website (www.attotech.com)
3 Select "Update"
4 Reboot when the update has completed

9.3 Verifying and Modifying NVRAM for the Celerity Host Adapter

The ATTO Configuration Tool will display the current parameters in NVRAM for each ATTO host adapter channel. Select the appropriate host adapter channel from the device tree on the left. This will display all available NVRAM parameters for that channel. You can change these settings and there are several options to save.

Note

The adapter settings are on a per channel basis and should only be changed by advanced users

Below are the available NVRAM parameters (See 9.3-1 on page 22):

Node Name
The Node WWN (World Wide Name) assigned to this channel of the adapter.

Port Name
The Node WWN (World Wide Name) assigned to this channel of the adapter.

Boot Driver
If enabled and the system supports EFI (Extensible Firmware Interface), you will be able to boot OS X off the Celerity Fibre Channel adapter. (Applicable to the 8 and 16Gb Celerity models)

Valid settings = (Enabled, Scan Only, Disabled) Default = Disabled

Note

Scan Only has no effect in Mac OS X.

Hard Address Enable
When a fibre channel loop is initialized, each device selects and configures itself to an available ID. Hard Address Enable permits the host to select a value entered in the Hard Address field.

Valid settings = (Enabled, Disabled) Default = Disabled

Note

It is recommended to leave this disabled.

Hard Address
The value used as the FC-AL hard address.

Valid settings = (0-125) Default = 0

Note

It is recommended not to change this setting.

Execution Throttle
Not used in OS X.

Frame Size
Changes the size of the FC packet of information being sent.

Valid settings = (512,1024,2048) Default = 2048

Note

It is not recommended to change this setting.

Device Discovery
Specifies the type of device discovery the adapter performs.

Valid settings = (Port WWN, Node WWN) Default = Port WWN

Note

This should not be changed in a multipathing environment.

Connection Mode
Arbitrated Loop (AL): Connects to either a Arbitrated Loop or a Fabric Loop Port (FL Port) on a switch. Point-to-Point (PTP): Connect to a direct Fabric connection, such as a F Port on a Fibre Channel Switch. AL Preferred : Allows the adapter to determine what kind of connection to use, but tries to connect in Loop Mode first, then Point-to-Point mode. PTP Preferred: Allows the adapter to determine what kind of connection to use, but tries to connect in Point-to-Point Mode first, then Loop mode.

Valid settings = (PTP Preferred, AL Preferred, PTP, AL) Default = PTP Preferred

Data Rate
Selects the Fibre Channel transmission rate. Auto indicates the adapter determines the rate based upon the devices connected.

Valid settings = (Auto, 16Gb/sec, 8Gb/sec, 4Gb/sec, 2Gb/sec, 1Gb/sec) Default = Auto

Interrupt Coalesce

Specifies the time period an adapter chip delays an interrupt allowing the adapter to queue interrupts.

Valid settings = (None, Low, Medium, High) Default = None

Note

Setting this value to high may be helpful when doing small IOs.

Port Down Retry Count

This variable determines the number of times the driver retries a command due to a port being logged out. This is a target initiated parameter. The device driver must think the target has logged out of the adapter. The count is decremented every time a command is aborted due to no login and retried when the device comes back. The Port Down Retry Count has nothing to do with time. This is a count based on each command sent from the adapter. This feature is designed for an erratic link, where a port would log out and in multiple times.

Valid settings = (0-255) Default = 8

Example 1: Port Down Retry Count = 3

- A command is sent out on the physical link..
- The target port logs out for 500msec.
- The target port logs back in. The command is retried (count is now 2).
- The target port logs out for 500msec.
- The target port logs back in. The command is retried (count is now 1)
- The target port logs out for 500msec.
- The target port logs back in. The command is retried (count is now 0)
- The target port logs out for 500msec.
- The error is now passed to the OS layer

Example 2: Port Down Retry Count = 3; Link Down Timeout = 10

- Same scenario as above. The Link Down Timeout does not matter since all retries will be exhausted before the 10 second Link Down Timeout expires.

Link Down Timeout

The amount of time (in seconds) the device driver waits for the link to be down before reporting an error to the OS and removing the device from the system.

Valid settings = (0-255) Default = 0 (Special case equivalent to 1 second)

This only comes into effect when there is one path left to a Target/LUN. The following conditions:

- The link from the HBA to the target or switch could go down.
- The target could also explicitly logout.
- The controller can fail or be taken offline for maintenance.

Example: Link Down Timeout =0 (special case = 1 second)

- If the adapter port senses that the link has gone down to the device, the driver will wait 1 second and then report and error to the OS.

Spinup Delay

Not used in OS X.

Load

This option allows the user to update the NVRAM parameters of a channel from a file.

Note

NVRAM settings are loaded from a file on a channel basis.

Save

This option allows the user to update the NVRAM parameters of a channel from a file.

Note

NVRAM settings are saved to a file on a channel basis.

Commit

This option allows the user to save the NVRAM parameters of the particular channel.

Note

NVRAM settings are saved on a channel basis.

Restore

This option allows the user to restore the NVRAM parameters to what is currently present in NVRAM, if changes have been made and not committed.

Note

NVRAM settings are restored on a channel basis.

Defaults

This option allows the user to restore the NVRAM parameters to defaults.

Note

NVRAM settings are defaulted on a channel basis.

Exhibit 9.3-1 NVRAM Settings

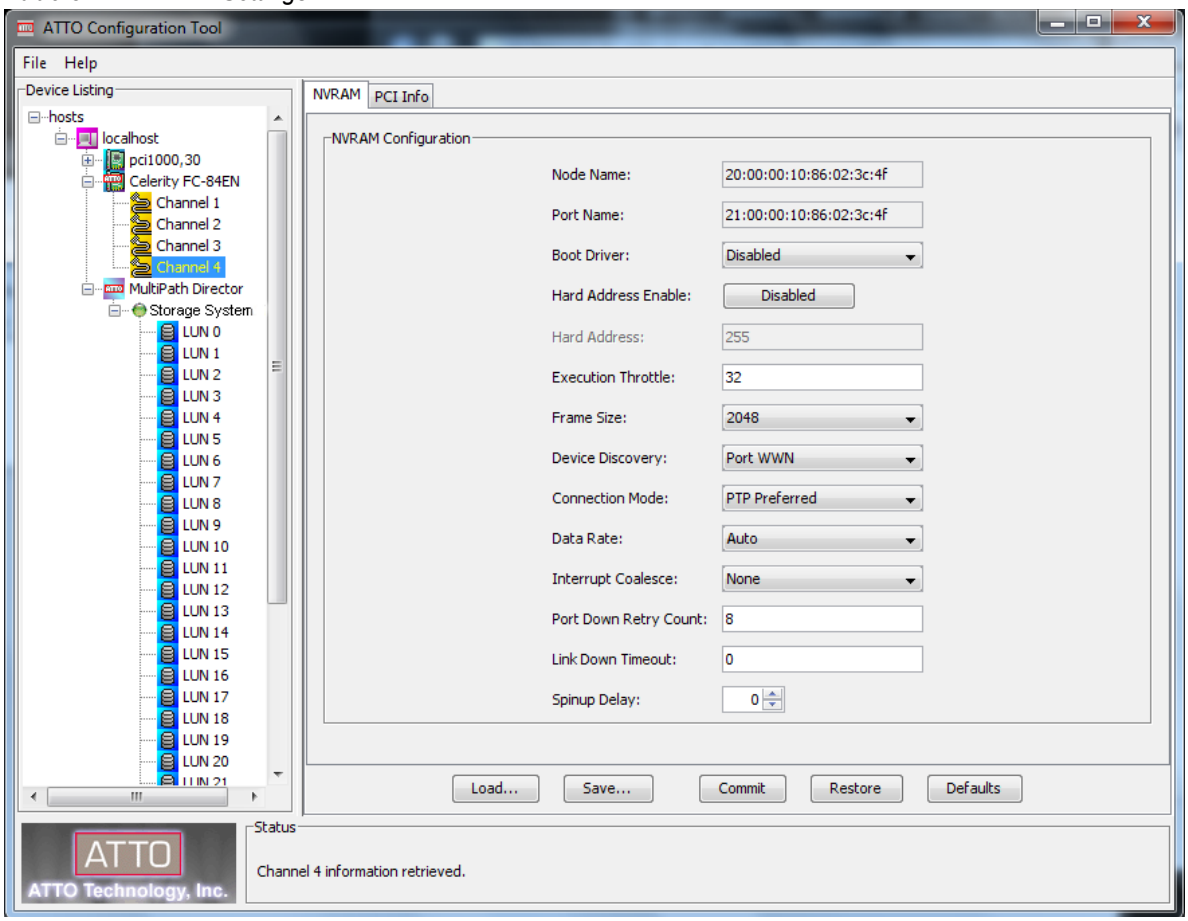

10 CLI Applications with Help and Usage Example

Atmpcfg.exe

This tool can be used to save the configuration settings for Celerity HBA's MultiPath Director.

Options:

-b {bus}
Select a bus number for the operation, startsat 0, default is all busses.

-c {channel}
Select a specific adapter channel for the operation, starts at 1, all channels are selected by default.

-h
Display extended help.

-i {infotype}
Print multipathing information. Information types are:

all *All information*

device *All device (target) information*

lun *All logical unit information*

-j
Reset path statistics.

-k
Locate a path (toggle on/off).

-l
List the multipathing directors in the system.

-m {mode}
Set the specified path mode. Modes are preferred, alternate, or disabled.

-n {policy}
Set the specified load balancing policy. Policies are pressure, roundrobin, or queuedepth.

-o {failoverMode}
default / fbdisabled.

-p {path}
Select a path for the operation, starts at 1, default is all paths. Use 0 to omit path information.

-r
Delete the saved multipathing configuration.

-s
Save the current multipathing configuration.

-t {id}
Select a device (target) ID for the operation, starts at 0, default is all devices.

-u {lun}
Select a logical unit for the operation, starts at 0, default is all logical units.

-v
Display non-error messages.

Atmpinfo.exe

This tool can be used to display information for the Celerity HBA's multiple paths to the storage.

Options:

-b {bus}
Select a bus number for the operation, starts at 0, default is all busses.

-c {channel}
Select a specific multipathing director for the operation, starts at 1, default is all channels.

-h
Display extended help.

-i {infotype}
Print multipathing information. Information types are:

all *All information*

device *All device (target) information*

lun *All logical unit information*

-l
List the multipathing directors in the system.

-p {path}
Select a path for the operation, starts at 1, default is all paths. Use 0 to omit path information.

-t {id}
Select a device (target) ID for the operation, starts at 0, default is all devices.

-u {lun}
Select a logical unit for the operation, starts at 0, default is all logical units.

-v
Display non-error messages.