

Altos S205F MPIO User Guide

Issue 1.0

1 November, 2004

- Acer and the Acer logo are registered trademarks of Acer Incorporated. Other company's product names or trademarks are used herein for identification purposes only and belong to their respective companies.
- All other names, brands, products or services are trademarks or registered trademarks of their respective companies.

Notices

Changes may be made periodically to the information in this publication without obligation to notify any person of such revision or changes. Such changes will be incorporated in new editions of this manual or supplementary documents and publications. This company makes no representations or warranties, either expressed or implied, with respect to the contents hereof and specifically disclaims the implied warranties of merchantability or fitness for a particular purpose.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopy, recording, or otherwise, without the prior written permission of Acer Incorporated.

Issue 1.0

1 November, 2004

Contents

| | |
|---|-----------|
| Preface | v |
| Revision History | vi |
| 1 MPIO Driver Definition and Installation | 1 |
| 1.1 MPIO Driver Definition | 1 |
| 1.2 MPIO Driver Installation Procedure | 2 |
| 1.3 Uninstalling the MPIO Driver | 2 |
| 2 Getting Started | 3 |
| 2.1 Multi-path Enabled Data Screen | 3 |
| 2.2 Advanced Configuration Screen | 6 |
| 2.2.1 <i>Retry Interval</i> | 6 |
| 2.2.2 <i>Retry Count</i> | 6 |
| 2.2.3 <i>Path Verify</i> | 7 |
| 2.2.4 <i>Path Verification Period</i> | 7 |
| 2.2.5 <i>PDO Remove</i> | 8 |
| 2.2.6 <i>Path Load Balancing</i> | 9 |
| 2.2.7 <i>Advanced Connection Host (debug option)</i> | 9 |
| 3 Configuration Options | 11 |
| 3.1 Dual Controller, Dual Ported Host, Dual FC Connection with MPIO | 11 |
| 3.2 Dual Controller, Dual Ported Host, Dual FC Connection via FC Switch with MPIO | 12 |
| 3.3 Advanced Configuration Example | 13 |
| 4 Frequently Asked Questions | 15 |

Preface

This guide describes the Multipath I/O (MPIO) driver for the Altos S205F FC to SATA RAID storage system, designed to operate as an integral part of the Windows Operating System. The need to have systems available continuously, 24 hours a day and 7 days a week, in the enterprise and electronic commerce space is no longer just a benefit but now a definite requirement. System availability is an important consideration for IT managers in determining the direction for hardware, storage and operating systems.

Storage media is a major point of failure in any server. In the past, several methods have been employed to reduce the probability of failure and to allow recovery after a storage failure has occurred. For example:

- Redundancy within the drives using RAID 5.

However, these solutions address only media problems. If a controller or the physical path to the media (for example, the Fibre Channel cable) fails, no amount of spindle redundancy will keep the machine up and running. The data on the media is still coherent, but there is no way to access it.

The MPIO driver provides the infrastructure to identify and eliminate the multiple single points of failure between the host bus and the physical storage media.

Chapter 1

MPIO Driver Definition and Installation

1.1 MPIO Driver Definition

The MPIO driver has been developed to enable the connection of a Altos S205F FC storage enclosure to a SATA RAID controller under a Windows Operating System without the need for a HBA vendor specific Failover driver. It includes the following features:

- Effectively HBA agnostic.
- Upper layer filter which sits over the standard HBA port driver.

The MPIO driver is integrated into the Windows Operating System rather than the HBA driver, hence the MPIO driver is agnostic of the HBA. Traditionally, the MPIO driver has been implemented by the HBA vendor and has been propriety to the HBA. For example, Qlogic released Sansurfer and QLDirect and Emulex Lightpulse.

The primary goals for the MPIO driver on the Altos S205F are:

- 1 Dynamic configuration of the multipath bus driver and disks based on installed adapters and devices, including boot and page file devices, and dynamic response to events occurring in the system.
- 2 Transparent system architecture that relies on Plug and Play, maintains compatibility with existing driver stacks, and does not break other technologies in the system.
- 3 Dynamic load balancing via multiple paths.
- 4 Enable control of path failover via exposing the path failover mechanism to the user to allow advanced configuration.

1.2 MPIO Driver Installation Procedure

Note Please make sure there are no other MPIO drivers installed on this server. Otherwise, a conflict could result. Uninstall old drivers if necessary.

The MPIO driver is installed by the means of the *setup.msi*, which guides the user through the following installation procedure:

- 1 Double click the **Setup application** file from Windows Explorer.
- 2 Follow the on screen prompts and install in the default directory.
- 3 A reboot of the Operating System may be required depending on the combination of Operating System and HBA. This will be indicated during the installation process

Note The MPIO driver is only supported on:

- Windows 2000
- Windows 2000 Advanced Server
- Windows XP
- Windows 2003 (all editions)

The MPIO driver is accessed via the Control Panel. The options in available are discussed in [Chapter 3 on page 11](#).

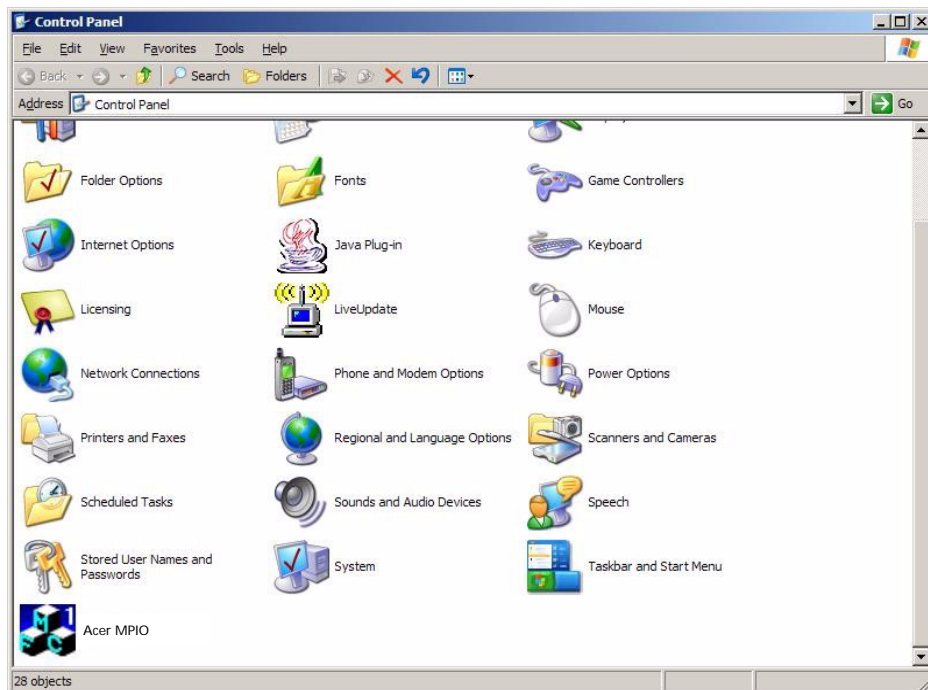


Figure 1–1 Control Panel

1.3 Uninstalling the MPIO Driver

The MPIO driver can only be uninstalled via the **Add/Remove Program** option in the Control Panel. This is the standard Windows removal procedure.

Chapter 2

Getting Started

The topology samples of this document are used to describe how MPIO works and its features. For formal solutions we recommend that users refer to "Altos S205F/S200F Product Configuration Guide" or contact Acer technicians for further details.

There are only two screens available in the MPIO driver, the Multi-path Enabled Data screen and the Advanced Configuration screen.

2.1 Multi-path Enabled Data Screen

A typical Multi-path screen is shown in [Figure 2-1](#), displaying all the available LUNs to the host system and the number of paths available.




Figure 2-1 Multi-path Data Screen

The LUN serial number is presented via the SCSI Inquiry pages 0x80 and 0x83.

Each device that the host can discover is queried using the **SCSI Inquiry** command to discover the multiple paths to a the same device. The LUNs are shown in the MPIO driver application by serial number.

Each LUN can be expanded to show the various paths to each LUN.

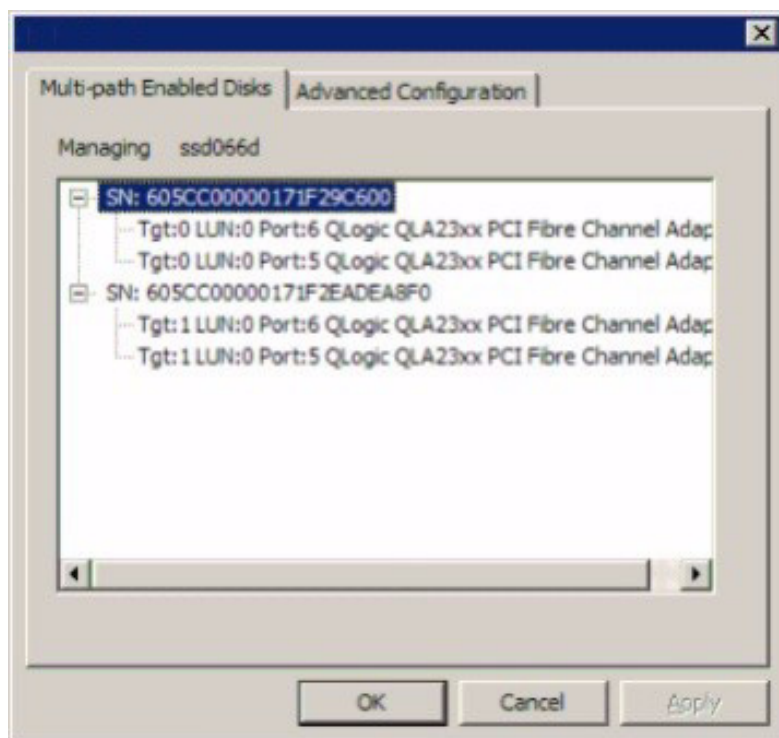


Figure 2–2 Available LUN Paths

The example shown in [Figure 2–2](#) is where there are two paths to each LUN. The configuration used to generate the two paths to each LUN is shown in [Figure 2–3](#). There are many different configurations which can be used to generate the same result i.e. multiple paths to the same LUN. These are described in [Chapter 3 on page 11](#).

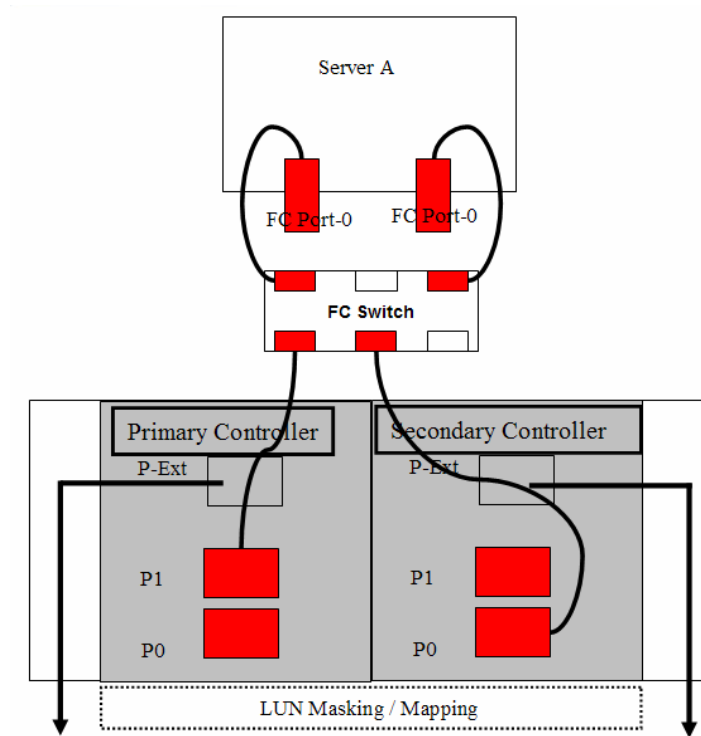


Figure 2–3 Multi-path Configuration

In order to achieve the multiple paths, the LUN for each controller was presented from:

- Host Port 1 (P1) on the primary controller
- Host Port 0 (P0) on the secondary controller

2.2 Advanced Configuration Screen

A typical Advanced Configuration screen is shown in [Figure 2–4](#).

The MPIO driver is installed with default settings to allow operation immediately after installation without any further setup. For the majority of Altos S205F installations the default settings are suitable and no further adjustments are required.

The Advanced settings can be found under the **Advanced Configuration** tab, shown in [Figure 2–4](#). The settings should only be altered after thorough testing of the new parameters has been carried out otherwise unpredictable results may occur.

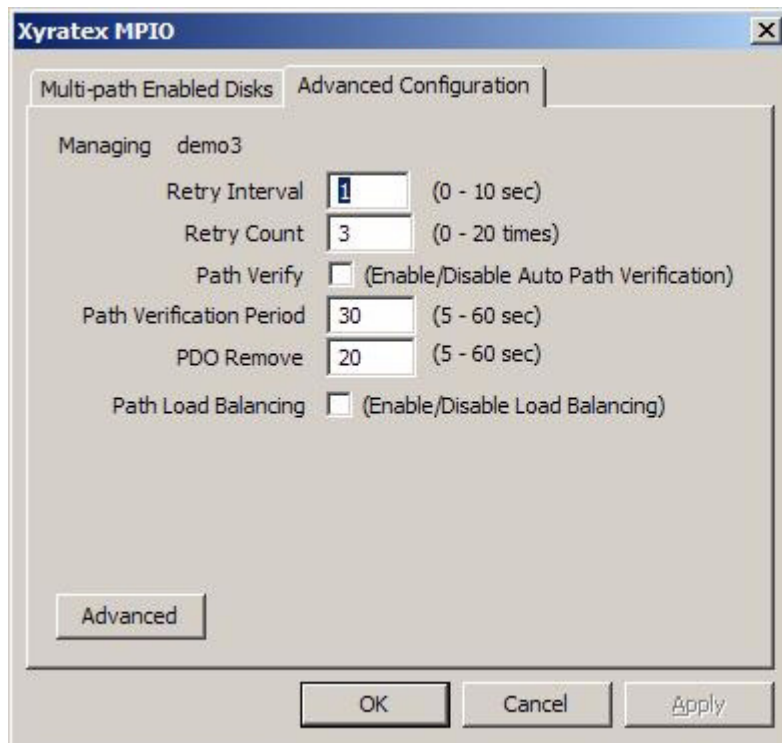


Figure 2–4 Advanced Configuration Screen

2.2.1 Retry Interval

The amount of delay (in seconds) between each retry of an I/O request. The interval has a maximum value of 10 seconds. If an I/O request fails it will be retried every X seconds as specified by this value.

2.2.2 Retry Count

The number of times to attempt the I/O request before failing the I/O request. Therefore the number of seconds before a path to a LUN will be failed is:

= Retry Interval X Retry Count

If the Retry Count is set to 0, then the Retry Interval is not applicable as the failed I/O request will not be retried. The path to the LUN will be immediately failed. This is not a recommended setting.

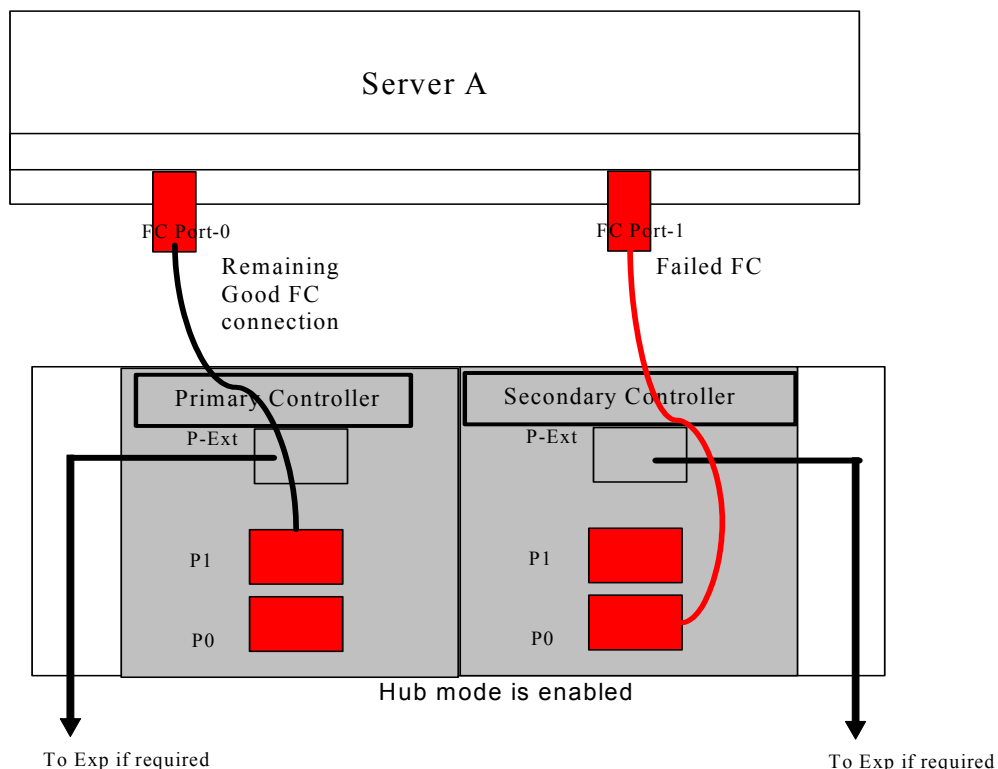


Figure 2–5 Retry Count Topology

The Retry Count and the Retry Interval determine how quickly a FC connection path is failed.

If an I/O fails, then it will be retried after the number of seconds specified by the Retry Interval value. The I/O will be retried for a maximum number of attempts as specified in by the Retry Count value. Once the maximum number of retries has been exceeded then the FC connection path is failed and the I/O will be moved to the remaining valid FC connections, refer to [Figure 2–5](#).

2.2.3 Path Verify

When working with multiple paths it is also often helpful to proactively determine the health of a path. This allows a non-operational path to be taken out of service prior to an I/O request being submitted to the path. MPIIO provides a means to verify each I/O path on a time interval to determine the health of the paths.

2.2.4 Path Verification Period

The time delay in seconds between attempts to verify each paths validity. The MPIIO driver sends a test unit ready I/O requests to verify the health of the path. If the I/O fails the path is reported as failed.

The Path Verification Period defines how often the path is verified using a SCSI Test Unit Ready command. If the **Path Verify** option is not selected then this value is ignored.




Figure 2–6 FC Connections

The topology in [Figure 2–6](#) shows two FC connections. If Auto Load Balancing is not selected, then only one of the FC connections is used for active I/O. The other FC connection is passive and will only be used if the active FC connection fails.

While the FC connection is passive, a **Path Verify** command (SCSI Test Unit Ready) is issued to ensure the connection is still valid.

The frequency of the **Path Verify** command is set via the Path Verification Period value

2.2.5 PDO Remove

Because there are circumstances where ALL paths may be unavailable for a short period of time there is a need to keep the MPIO virtual PDO from being eliminated until at least one path recovers.

This is the timeout value used by the MPIO driver to wait until an available FC connection is available, after all FC connections have disappeared.

2.2.6 Path Load Balancing

Path Load Balancing allows the I/O to be balanced through each access path so that all paths are utilized equally.

Path Load Balancing should be used for the majority of applications as it utilizes all available FC connections via Active/Active connections. Normally 99.99% of time, the host will have all FC connection paths available.

If Path Load Balancing is not used, then one FC connection path is active and all other paths are passive. A passive FC connection path will only be used if the active FC connection path is failed. Only applications which do not have high performance requirements should be considered for Non Path Load Balancing.




Figure 2–7 Path Load Balancing Topology

If Path Load Balancing is utilized in this example, the maximum performance is 400MB/s (2 x 200MB/s)

If Path Load Balancing is not utilized in this example, then the maximum performance is 200MB/s (1 x 200MB/s).

2.2.7 Advanced Connection Host (debug option)

The option Advanced Connection Host (debug option) under the Advanced Configuration tab is intended for development purposes only.

Chapter 3

Configuration Options

This chapter describes configurations and topologies which can be used to achieve multiple paths to the Altos S205F LUNs. The following sections only contain a subset of all the possible configurations and are by no means exhaustive.

3.1 Dual Controller, Dual Ported Host, Dual FC Connection with MPIO

| | |
|--|---|
| Cabling | Connection to p0 on secondary and p1 on primary |
| LUN Mapping | To both ports |
| Altos S205F InterHub Mode | On |
| Multipathing | Yes |
| LUNs Available | 32 per port |
| Number of LUN 0 available to host | 2 (Dual pathed) |

The configuration shown in [Figure 3–1](#) uses the internal hub on the Altos S205F to achieve the multiple paths to each of the LUNs. In order for this to occur the LUNs from each controller need to be presented via both ports.

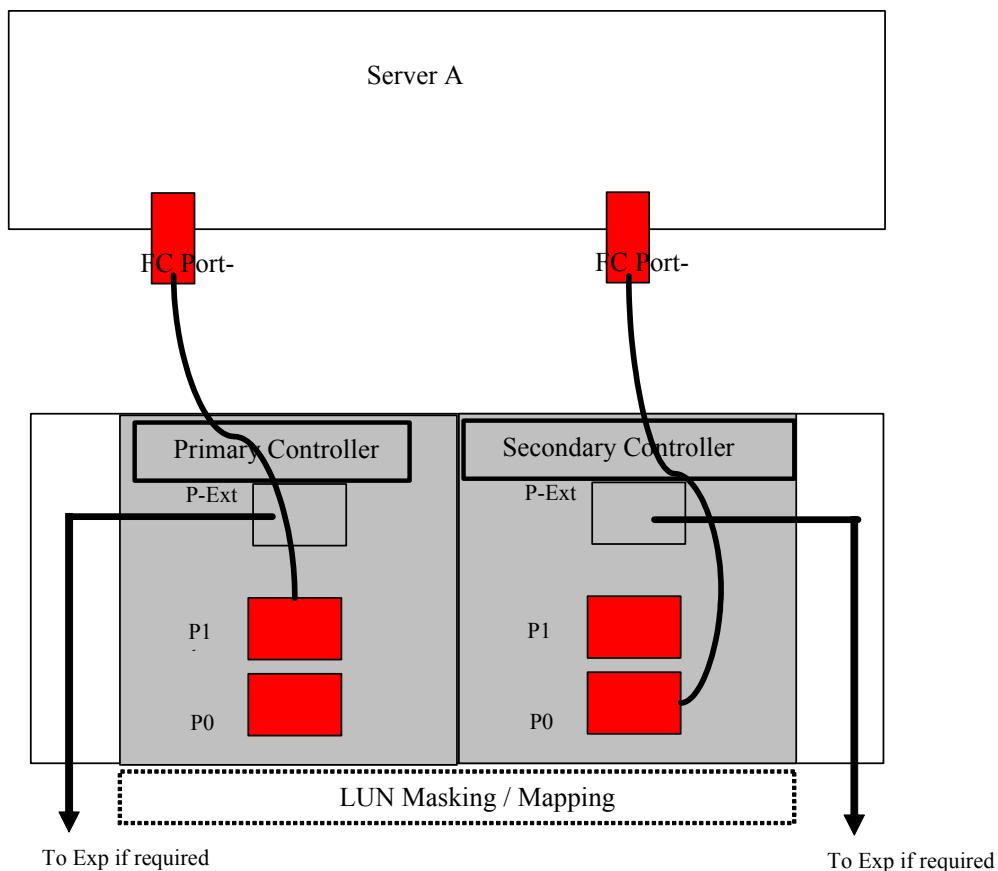


Figure 3–1 Dual Controller, Dual ported Host, Dual FC Connection Configuration

3.2 Dual Controller, Dual Ported Host, Dual FC Connection via FC Switch with MPIO

| | |
|--|---|
| Cabling | Connection to p0 on secondary and p1 on primary |
| LUN Mapping | To both ports |
| Altos S205F InterHub Mode | Off |
| Multipathing | Yes |
| LUNs Available | 32 per port |
| Number of LUN 0 available to host | 2 (Dual pathed) |

The configuration shown in [Figure 3–2](#) uses the Fibre Channel switch connected to the Altos S205F to achieve the multiple paths to each of the LUNs. In order for this to occur the LUNs from each controller need to be presented via both ports. The Altos S205F internal hub must also be off.

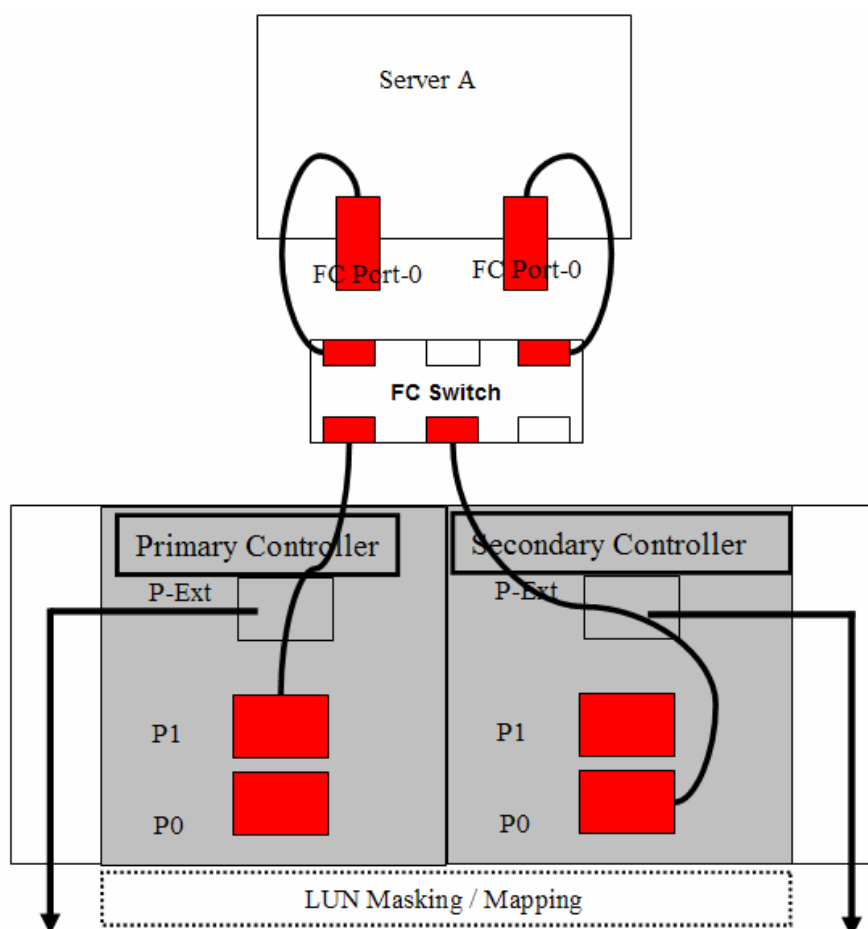


Figure 3–2 Dual Controller, Dual ported Host, Dual FC Connection via FC Switch Configuration

3.3 Advanced Configuration Example

As mentioned in section 2.2 on page 6, the default settings of the MPIO driver should allow operation without any alteration in the majority of applications. This section outlines an example where the Advanced Configuration settings may be changed.

If the host application is time sensitive, for example, a transactional based database, then it may not be applicable to retry the I/O multiple times with a large period of time in between the retries.

In this scenario, it may be applicable to alter the Retry Count and Retry Interval to minimum values as response time is more critical than multiple FC connection paths.

If this was done then a FC connection path failover is more likely as the host (MPIO driver in particular) is less tolerant of failed I/O commands.

The Retry Interval and Retry Count should only be altered after the “new” values have been tested first. Otherwise these “new” values may have an undesired effect on the host system.

Chapter 4

Frequently Asked Questions

Q: What happens if the “Hub Mode” on the Altos S205F is used and how can multiple paths be created?

A: Using the Hub Mode will allow a single Fibre Channel link to discover the LUNs on both Altos S205F controllers. The LUNs for each controller must be presented through both ports on each controller.

Q: Can the MPIO driver be used with a single Altos S205F?

A: The MPIO driver can be used on a single Altos S205F by using two connections into the Altos S205F controller. Please be aware there is still a single point of failure as there is only a single controller present.

Q: How can multiple paths be created without a FC switch or hub present?

A: As explained in [Chapter 3](#), the Altos S205F controller have an internal hub present which can be enabled. This will allow multiple paths to each LUN.

Q: Is it possible to generate more than two paths to the same LUN?

A: Yes, using a Fibre Channel switch this is possible. The number of HBA connections and/or connections to the Altos S205F controllers would need to be increased.

Q: Does the MPIO driver operate on any other storage device?

A: No, the MPIO driver will only operate with the Altos S205F FC to SATA product.

Q: Is there a version of the MPIO device driver available for the Linux Operating System?

A: Unfortunately, due to the MPIO driver being integrated into the Windows Operating System there is not a version for the Linux Operating System. User may use multi-path software or a driver provided by FC HBA or other 3rd party vendor.

Q: Is it possible to connect to more than one Altos S205F from the same host?

A: Yes, it is possible but the multiple Altos S205F controllers would have to be connected via the a Fibre Channel switch to provide enough connections.

Q: Does the MPIO driver operate with a Altos S200F directly connected to the host?

A: No, the MPIO driver is only designed for the Altos S205F. Altos S200F enclosures can only be supported when operating behind the Altos S205F FC to SATA RAID controllers.

Q: Is there any graphic view of the connections between the host and the Altos S205F?

A: No, unfortunately the first release of the MPIO driver does not display any graphical connections between LUNs and hosts.