



Integrating Xyratex Storage with Apple Mac OSX

By Darren Attree



Notices

The information in this document is subject to change without notice.

While every effort has been made to ensure that all information in this document is accurate, Xyratex accepts no liability for any errors that may arise.

© 2008 Xyratex (the trading name of Xyratex Technology Limited). Registered Office: Langstone Road, Havant, Hampshire, PO9 1SA, England. Registered number 03134912.

No part of this document may be transmitted or copied in any form, or by any means, for any purpose, without the written permission of Xyratex.

Xyratex is a trademark of Xyratex Technology Limited. All other brand and product names are registered marks of their respective proprietors.

Issue 1.6 | November, 2008

Contents

Introduction	2
Hardware	2
Configuration of ATTO HBA	2
Configuration Procedure of Apple (LSI) HBA.....	3
F54xxE Setup.....	3
Host Port Configuration	3
FC Connectivity.....	3
SAN Mappings with ATTO HBA	3
SAN Mappings with LSI HBA.....	3
Node Names	3
Spoof Cache Sync Flush	3
2TB LUN Limit Removal	4
Configuration of ATTO Multipath Driver	4
Display of LUNs and Mounted Drives within OS.....	6
I/O Testing of Mounted Drives	7
AJA Video Performance Testing of Mounted Drives.....	7
FC Switch Configuration.....	8

Introduction

This report confirms compatibility between the F5402E, the F5412E and the F5404E with Mac OSX V10.5.3. This compatibility was confirmed on a test bed including ATTO® Celerity and Apple® HBAs and Brocade® SilkWorm and QLogic® 5602 switches. Details regarding models, revisions, test procedure, and configuration are contained in this report.

The purpose of this document is to describe the interoperability and best practices between the Xyratex FC to SAS/SATA storage products, and the Apple Mac server. This will include any known limitations.

Hardware

- Mac OSX V10.5.3
 - 2X 3GHz Dual Core Intel® processors
 - 2G RAM
- ATTO FC42ES dual ported HBA — PCIe
 - 2X 3GHz Dual Core Intel processors
 - 2G RAM
- Apple (LSI7204EP) dual ported HBA PCIe
- F5402E / F5412E with 2.4 FW bundle
- F5404E with 2.3 and 2.4 FW bundles
- Xyratex Qualified SAS or SATA drives
- Brocade SilkWorm 200E 16 port 4G FC switch
- QLogic 5602 16 port 4G FC switch
- XSAN version 2.1.1

Configuration of ATTO HBA

- Install ATTO Configuration Tool version 'osx_app_configtool_320.dmg'
- Upgrade HBA FW to version 'osx_fsh_celerity_071113.dmg'
- Install HBA Multipath Driver version 'osx_drv_celerity_325MP.dmg'
- Speed — Auto
- Topology — Auto
- Direct connect works fine
- Connected via Brocade 200E or QLogic 5602 works fine

Note: Without installing the ATTO Configuration Tool and upgrading the HBA FW, it is not possible to install the driver and no HBA HW will be detected.



Configuration of Apple (LSI®) HBA

- Speed — Auto
- Topology — Auto
- FW version 1.3.20.0
- Direct connect with a single host connection to the storage works fine. HA does not.
- Connected via QLogic 5602 with a single host connection to the storage works fine. HA does not.

Note: LSI does not supply a Multipathing Driver. As such when using an LSI HBA, the Mac Multipathing driver must be used.

F54xxE Setup

Host Port Configuration

- Speed — Auto
- Topology — Auto

FC Connectivity

- C0P0 to HBA Port 0
- C1P1 to HBA Port 1

SAN Mappings with ATTO HBA

- All LUNs mapped to HBA Port 0 WWPN via C0P0
- All LUNs mapped to HBA Port 1 WWPN via C1P1
- Processor LUN disabled

SAN Mappings with LSI HBA

- All LUNs mapped to HBA Port 0 WWPN via C0P0
- Processor LUN disabled

Node Names

- N/A

Spoof Cache Sync Flush

- Enabled

Note: The LSI HBA is not supported within an HA environment (dual RAID controllers connected to multiple HBAs within a single host to provide redundancy).

2TB LUN Limit Removal

- Connect RS232 cable to C0
- 115200,8,1,N,1
- Ctrl E
- Enter the 'Configuration' menu
- Type 'enable advanced configuration 4' and press enter
- Enter the 'Advanced Configuration' menu
- Enter the 'Configure Advanced Performance Options' menu
- Enter the 'Configure INQUIRY SPC Fields' menu
- Set to SPC-3
- Exit and save settings

Note: 2.4 FW onwards does not require SPC-3 to be manually configured as it is ON by default.

Configuration of ATTO Multipath Driver

The ATTO Configuration Tool can be used to view and configure LUN multipathing for devices attached to ATTO Celerity HBAs.






Click on the device in the device tree, and then click on the Paths tab on the right to display the multipathing window.

The Paths tab displays path information for the device. The upper half of the window displays information for all paths:

- Status — The overall status of all paths
- Online — All paths are functioning normally
- Degraded — At least one, but not all paths have failed
- Offline — All paths have failed
- Changed — New paths are present for a previously configured device. You should reconfigure the device to include the new paths
- Load Balance

Transfer Count — The total number of bytes transferred on all paths.

Paths — A list of all paths to the device (see Figure 1). If you clicked on a LUN in the left navigation pane, the icon for each table entry indicates the state of the path:

-  Active — The path is being used for I/O (Active Path)
-  Inactive — The path is connected, but not being used for I/O (Passive Path)
-  Failed — The path is not connected
-  Disabled — The path is connected, but has been disabled
-  A single path is connected to your device

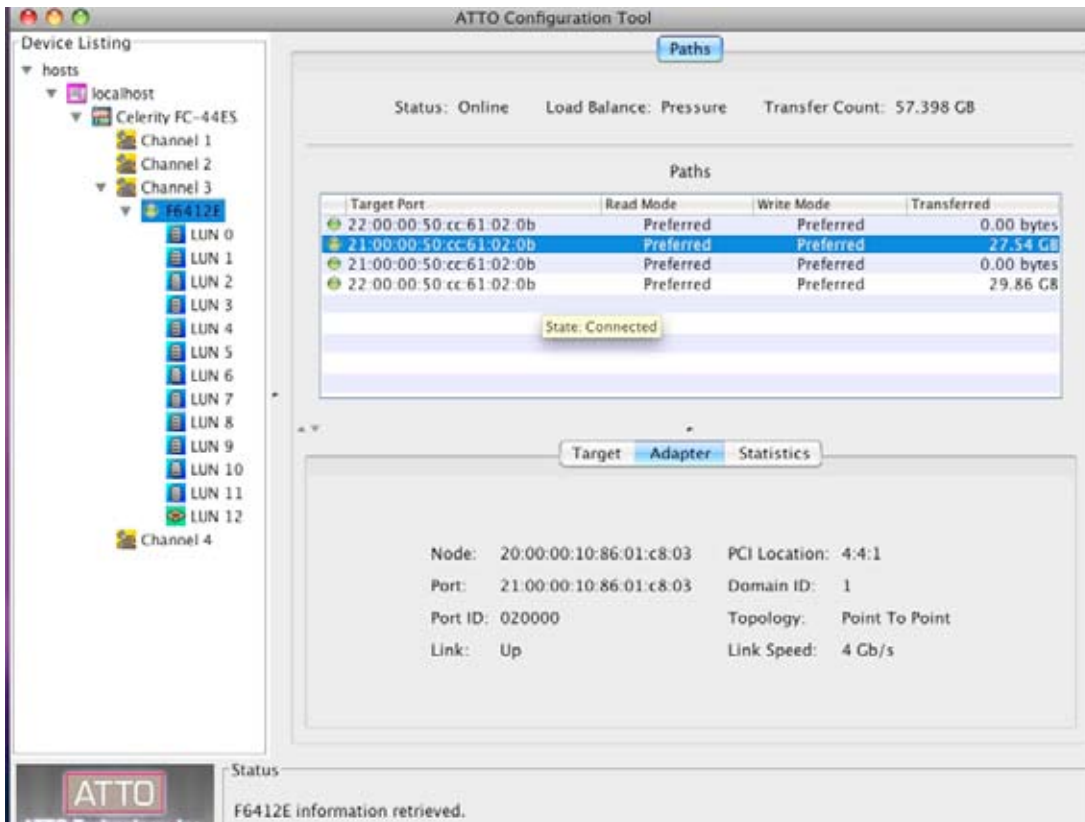


Figure 1

Important Note: Even though a dual ported HBA is used and the same LUNs are presented to both ports (Channels). The multiple paths to the LUNs are only shown on one channel. e.g. LUN 1 above is only seen on channel 2. However there are 4 paths to it via a FC switch and the 2 HBA channels (it just looks like only 1 channel is physically connected).

- Right click on either the LUN or the Channel you wish to configure multipathing. Select 'Setup' — A configuration wizard will appear.
- Select the load balancing policy for all paths to the device:
 - Pressure — The path with the fewest bytes being transferred is selected for I/O
 - Queue Depth — The path with the fewest commands being processed is selected for I/O
 - Round Robin — The least used path since the current configuration was setup is selected for I/O

Note: The above are not used for Xyratex and 'Pressure' is the option to use as the following setting that has to be used for Xyratex force an active and passive path rather than load balancing.

- For each path, select how each path is used for reading and writing data
 - Primary — The path will be used to transfer data
 - Alternate — The path will be used to transfer data if a primary path is not available
 - Disabled — The path is not used to transfer data

Note: A primary and alternate path must be configured for all LUNs. DO NOT configure multiple primary paths to the same LUN as this will cause load balancing to be used which will reduce performance.

- Click 'Next' until you have configured all paths
- Click 'Finish' and then 'Save' all settings. The new settings will take effect immediately

Note: Once you have configured the device, the new settings will persist across system reboots. Even if you do not change the configuration for a path, its configuration will be saved. Paths that are saved are viewable in the Configuration Tool when not connected.

Display of LUNs and Mounted Drives within OS

- Open the Mac 'System Profiler'
- Expand the Fibre Channel tab
 - The user will then be presented with the 2 FC ports of the LSI HBA and the LUNs presented from the F54xxE under both ports (i.e. the same LUNs twice and Round Robin load balance's them which severely reduces the RAID performance).

Note: The above would be an unsupported configuration due to the very poor performance in a HA environment.

- The user will then be presented with the 2 FC ports of the LSI HBA and the LUNs presented from the F54xxE under both ports (i.e. the same LUNs twice and Round Robin load balance's them which severely reduces the RAID performance).

Note: OSX has a Multipathing driver installed by default. This cannot be disabled!

- Open the Mac 'Disk Utility'
- From this screen you will see only one instance of each of the LUNs displayed as a Drive.
- From the 'Disk Utility' page you are able to create either individual 'Volumes Mount Points' within the OS by using the 'Erase' button. Or create a basic OS SW RAID set under the 'RAID' tab.
- The typical mount points for basic drives are:
`/volumes/{Name_You_Have_Specified}`

I/O Testing of Mounted Drives

As OSX is based on Unix the user is able to use the basic Unix dd command to perform I/O tests.

- Open a 'Basic' Unix terminal screen
- Type the command:

```
dd if=/dev/zero of=/Volumes/{Name you have specified} bs={Block size you wish to transfer seeing that each block is 512 bytes, so a 2MB transfer would require 4000} count={The number of files you wish to transfer — Typically use a value depending upon how large each transfer is (e.g. bs=2MBxcount=1000=2GB space required on disk)}
```

Overall disk performance can be measured within the Mac 'Activity Monitor'

- Only the following can be measured:
 - Overall reads
 - Overall writes
 - IOPs
 - BW
- The user is unable to differentiate between Sequential or Random

The performance monitor within StorView™ provides better granularity.

Note: Multiple terminal screens can be opened allowing multiple dd commands to be run to multiple LUNs simultaneously.

AJA® Video Performance Testing of Mounted Drives

The AJA Kona video test tool is the choice of performance bench marking for video editing using a Mac.

- Streaming write performance of user defined files and resolutions (e.g. 1080P, 1080i) to ensure adequate frame rate is attainable.
- Streaming Read performance of user defined files at defined resolutions (1080P, 1080i) to ensure adequate frame rate is attainable.

The F6412E storage with both SAS and SATA drives configured within RAID-5 arrays was able to attain the desired performance even at the highest resolutions.

FC Switch Configuration

To date only 2 FC switches have been tested, and is thus supported by Xyratex within a Mac to Xyratex FC to SAS/SATA system. Interoperability with other FC switches is currently being investigated with a view to gain support.

- Brocade SilkWorm 200E 16 port 4G FC Switch
 - FW V5.2.0.a
- All LUNs mapped to HBA Port 0 WWPN via C0P0
- All LUNs mapped to HBA Port 1 WWPN via C1P1
- Processor LUN disabled
- P0 to P4 marked as Storage Ports within the FC switch
- P5 to P8 marked as Host Ports within the FC switch
- C0P0 connected to P0 of FC switch
- C1P1 connected to P1 of FC switch
- P5 of FC switch connected to HBA Port 0
- P6 of FC switch connected to HBA Port 1
- All HBA ports and RAID controller ports within the same zone

Note: With Apple (LSI) HBA installed, the 'Brocade 200E' FC Switch is unable to see the HBA ports. As such the user cannot configure zones which the brocade requires to operate. The ATTO HBA is displayed correctly and as such works fine.

- QLogic 5602 16 port 4G FC Switch
 - FW V6.8.0.3.0
- All LUNs mapped to HBA Port 0 WWPN via C0P0
- All LUNs mapped to HBA Port 1 WWPN via C1P1
- Processor LUN disabled
- P0 to P4 marked as Storage Ports within the FC switch
- P5 to P8 marked as Host Ports within the FC switch
- C0P0 connected to P0 of FC switch
- C1P1 connected to P1 of FC switch
- P5 of FC switch connected to HBA Port 0
- P6 of FC switch connected to HBA Port 1
- Either no Zone or all HBA ports and RAID controller ports within the same zone

Note: Testing has confirmed that with either the Apple (LSI) HBA or ATTO HBA connected to the QLogic 5602 FC switch, ports will display correctly.

Note: Xyratex has concluded the Brocade SilkWorm xxx 2G xxxxxx does not work, and is no configuration containing this switch is supported.

About Xyratex

Xyratex is the ultimate partner to the storage industry. We are a leading provider of enterprise-class data storage subsystems and storage infrastructure manufacturing equipment & automation solutions. Working with over 50 A-list companies, Xyratex ships over 14% of the world's external storage capacity, and 75% of all 3.5" drives are processed using Xyratex test systems. With unmatched expertise and a history of innovation and technological excellence, Xyratex delivers products which are high-performance, energy-efficient and extremely reliable.

For more information, please visit www.xyratex.com

Xyratex Headquarters

Langstone Road
Havant
Hampshire PO9 1SA
United Kingdom

UK HQ

T +44 (0)23 9249 6000
F +44 (0)23 9245 3654

www.xyratex.com

Principal US Office

2031 Concourse Drive
San Jose, CA 95131
USA

USA Sales & Support

T +1 877 997 2839
T +1 877 XYRATEX



ISO 14001: 2004 Cert. No. EMS91560

©2008 Xyratex (The trading name of Xyratex Technology Limited). Registered in England & Wales. Company no: 03134912. Registered Office: Langstone Road, Havant, Hampshire PO9 1SA, England. The information given in this brochure is for marketing purposes and is not intended to be a specification nor to provide the basis for a warranty. The products and their details are subject to change. For a detailed specification or if you need to meet a specific requirement please contact Xyratex: www.xyratex.com.

x y r a t e x