



IBM System Storage™ Digital Media Storage Solution Installation Guide

***Selling and Deploying DS3000/DS4000/DS5000
in Mac, Linux, Windows Client, and StorNext Environments***

Con Rice
Modular Storage Product Specialist
LSI Corporation
12007 Sunrise Valley Drive, Suite 325
Reston, VA20191
Cell: 703-867-0012
Tel: 703-262-5418
e-mail:con.rice@lsi.com

Digital Media Storage Solution Installation Guide
October 2010
© Copyright 2010, IBM Corporation

Contents

INTRODUCTION	4
Mac, Windows and Linux Client Connectivity.....	5
StorNext File Sharing Environments	5
CURRENT SOLUTIONS	6
Supported Hardware Environments	6
Supported File Sharing Environments.....	7
Note: A customer seeking connection that has already deployed Apple Xsan MDC will need to transition to a Quantum StorNext MDC running either on Linux or Windows Server. However the Apple Xsan clients can remain.	8
SOLUTION COMPONENT DESCRIPTIONS	8
SAN Architecture criteria.....	8
Apple Clients.....	8
Windows Clients	8
Windows Server MDCs (Future Support).....	8
Linux StorNext Clients and MDCs.....	9
IBM DS Storage	9
Solutions.....	12
Apple Hardware Connect Only.....	12
Apple Servers and/or clients with no shared file access.....	12
Apple Homogenous Shared File	12
Apple Xsan Clients with Linux MDC (Meta Data Controller).....	12
Windows Client Hardware Connect Only	12
Windows clients with no shared file access.....	13
Linux Homogenous	13
Linux StorNext Clients with Linux MDC	13
Heterogeneous	13
Apple Xsan Clients, Linux StorNext Clients, Windows Clients and Linux MDC.....	13
APPLE, WINDOWS, and STORNEXT SUPPORT MATRIX.....	13
SOLUTION RESTRICTIONS AND RECOMMENDATIONS	15
RECOMMENDED RAID CONFIGURATIONS.....	16
SOLUTION INSTALLATION.....	17
QUANTUM / APPLE JOINT PRODUCT SUPPORT PROGRAM	20
Overview	20
Xsan Controlled SAN File System.....	20

StorNext Controlled SAN File System	20
Cooperation Effort Between Apple and Quantum.....	20
Closing a Case	20
MIGRATING FROM XSAN MDC TO STORNEXT MDC.....	21
ATTO Technology Inc.....	22
APPENDIX A: EXAMPLE CONFIGURATIONS.....	24
Large Heterogeneous StorNext configuration	24
Mid to large -sized “MAC Homogenous”StorNext configuration.....	26
Mid-sized “MAC Homogenous”StorNext configuration.....	27
Small “MAC Homogenous”StorNext configuration	28
Mid to large -sized Heterogeneous StorNext configuration.....	29
Additional Reference Links	31
Notices.....	31
Trademarks	32

Acknowledgements: Many people contributed to this document including:

Ryan A. French, Application Engineer
 ATTO Technology, Inc. · 155 CrossPoint Parkway · Amherst, New York 14068
 PHONE: 716-691-1999 ext. 188 - FAX: 716-691-9353
 EMAIL: rfrench@attotech.com
www.attotech.com

Mike Tenhulzen
 Alliance Marketing Manager
 LSI Corporation

Will McGrath - Software Alliance Manager
 Quantum Corp
 phone: 603.247.0350
 email: will.mcgrath@quantum.com

James Latham
 DS3000/DS4000/DS5000 Product Specialist
 LSI Corporation
 T 248-561-7445
 Email: jim.latham@lsi.com
www.lsi.com

INTRODUCTION

This paper outlines support of Windows Client (XP, Vista, and Windows 7), Linux (RHEL and SUSE) and Mac (OS X 10.5 – 10.6.x) with IBM DS3000, DS4000, and DS5000 Storage Servers. It also discusses support in StorNext environments. The solutions discussed arose because of customer demand seen in the media and entertainment industry and within the oil and gas industry.

Many media applications include a mixed environment with Windows clients, Mac clients, and sometimes even Linux clients requiring access to the same file system. Various software packages along the production and distribution chain work best on these disparate platforms. In the oil and gas industry this variety is not as prevalent, and Linux is the dominate operating system. Though even here, some amount of Windows systems may be encountered.

There are two basic aspects to these solutions: One is basic hardware connectivity, the second is the support for file sharing. It's important not to confuse the two. Connectivity to many Linux hosts is well established and outlined in IBM's on-line compatibility tool. Basic hardware connectivity to Apple and/or Windows Client hosts is now also outlined on IBM's on-line compatibility tool. The second aspect is File sharing with Quantum's StorNext software. This is also supported, but because of the added complexity requires very specific configurations.

The storage offerings in Apple environments in the past have been poor. The now discontinued Apple X-RAID storage products were never very highly regarded and neither is the current Promise storage. There has long been a desire to have simple and supported deployable solutions with StorNext and IBM DS Series Storage. In particular, there has been a strong demand for support of Apple servers and work stations with IBM DS series storage. IBM DS series now has that capability, through a very specific set of known, quality configurations.

In the Windows realm, IBM Storage has long supported Windows servers. In fact, SQL, Exchange, and many other applications have been a mainstay of our business. However, this support did not extend down to Windows clients like XP, Vista, and Windows 7. Windows clients, it was believed just didn't need access to the kind of high performing and reliable storage IBM offers. This has now changed. Windows clients are an integral part of the workflow processes within many industries, and support for them was clearly required.

The StorNext File system is a popular "shared" file system that is used for many publishing, oil & gas industry, and media applications. These applications operate in a variety of operating systems environments. Key to recognizing these opportunities is an understanding of the customer environment. Customers will often describe their production systems as a "workflow". The idea is to use high speed shared access to files to impose various processes. Using StorNext (or in the case of Apple, XSan) allows customers to achieve a high degree of efficiency in collaborative workflows allowing shared access to a virtual pool of storage

Mac, Windows and Linux Client Connectivity

In recent years, the workflow requirements of the video professional is driving the need for more sophisticated, high-end configurations, which require multiple users to share access to large amounts of data while maintaining the high-throughput performance required for the rich media and content applications they serve.

To support these solutions, a SAN-based storage subsystem is needed which supports Mac OS X, Windows Clients such as Windows 7 and Vista and sometimes Linux servers. The solution also needs to facilitate the sharing of common volumes in a resilient (failover and fallback) manner. Recently, numerous customers have asked IBM to provide a mid-range option to meet these needs.

The IBM Digital Media Solution combines high-performance IBM System Storage DS3000, DS4000 and DS5000 arrays used in conjunction with ATTO's Celerity line of quad/dual/single port 8Gb FC HBAs and Quantum's StorNext File System software to provide the infrastructure needed for content developers and video professionals to easily access, transfer, archive data, with the flexibility to scale performance and capacity to the various workflows.

The new IBM Digital Media solution addresses two common challenges in many of video streaming environments; the need to support Mac OS X, Windows Clients, and the need to provide high-availability storage access.

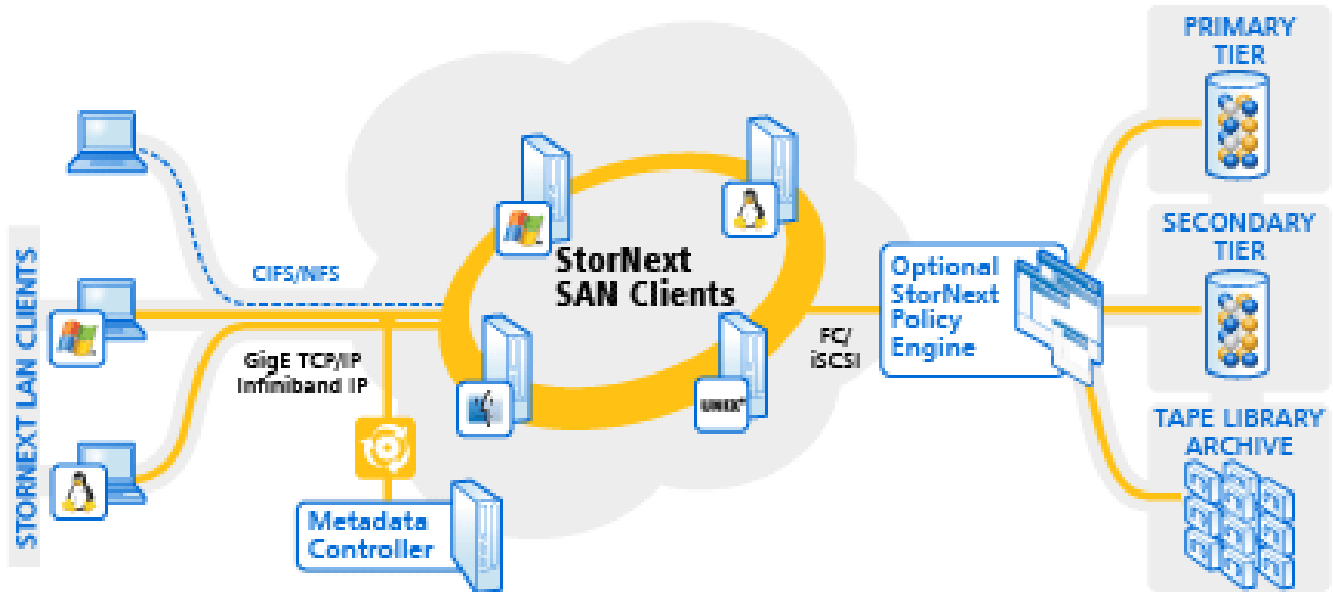
StorNext File Sharing Environments

Homogenous Linux installations often occur in the Oil and Gas industry for seismic processing and in some scientific computing environments. These environments generally include a pair of Linux servers acting as Meta Data Controllers using Quantum StorNext MDC software and a number of Linux servers and/or work stations using StorNext client software. Clients access files via Fibre channel SAN connections while reading and ask the MDCs for exclusive write control to files while writing.

An alternative configuration is the "storage node" or "Distributed LAN Client" (DLC) configuration. In this configuration, The MDCs maintain control, but LAN Servers act as Gateways to access the information and serve it to StorNext DLC clients via LAN connections. The LAN Gateway servers act as a storage node directly connected to the Fibre Channel or iSCSI DS Series storage and respond to clients over a high speed LAN connection with file access over Ethernet or Infiniband over IP using a high throughput StorNext data transfer protocol. In this case, there will be Quantum StorNext MDC licenses along with Gateway nodes and LAN Client licenses serving in a distributed (and shared) file system deployment.

In publishing and media environments, many applications require support for Apple MacOS. Apple installations typically use a shared file system called Xsan. Xsan enables multiple Mac OSX clients to connect to a single pool of storage and is fully compatible with Quantum's StorNext File System. StorNext is very popular in the media market because it allows shared block-level file system access while offering integrated archive management. However, Apple offers very limited support outside very specific configurations – with Promise as the RAID provider. A 100% Apple configuration must be used.

There are three components to the basic StorNext solution: Clients, Meta Data Controllers (MDCs), and the storage. There are also ancillary components such as the SAN, the TCP/IP network, and optional archive storage such as tape libraries or second tier disks. However, for selling and support purposes the components are software, support, and implementations services.



StorNext Environment

The announced support for the IBM Digital Media Storage Solution is the result of efforts between ATTO (HBAs), Quantum (StorNext), LSI (Storage Subsystems), and IBM. We have identified a set of hardware, firmware, and configuration parameters that will be supported. The set of configurations will be expanded in a logical manner over the next few months as additional choices are available and tested.

CURRENT SOLUTIONS

Supported Hardware Environments

	HBAs	Multipathing Software
Mac O/S	ATTO Celerity <ul style="list-style-type: none"> • FC-84EN, Quad Port • FC-82EN, Dual Port • FC-81EN, Single Port 	ATTO MultiPath Director
Windows 7, Vista, XP	ATTO Celerity <ul style="list-style-type: none"> • FC-84EN, Quad Port • FC-82EN, Dual Port 	ATTO MultiPath Director

	<ul style="list-style-type: none"> FC-81EN, Single Port 	
Windows Servers	ATTO (future support), see current compatibility matrix for alternative solutions	
Linux Servers	<p>In mixed MacOS and/or Windows Client (XP, Vista, Windows 7) environments: ATTO Celerity</p> <ul style="list-style-type: none"> FC-84EN, Quad Port FC-82EN, Dual Port FC-81EN, Single Port <p>In homogenous Linux environments, any currently supported HBA and failover:</p> <ul style="list-style-type: none"> QLogics, Emulex, and RDAC or MPP multipath as appropriate 	<p>In mixed, environments (those that include MacOS or Windows : ATTO MultiPath Director</p> <p>In homogenous Linux environments: RDAC or MPP as appropriate for the particular Linux implemented</p>

Using this table, support can be obtained for IBM DS3400, DS3500 (FibreChannel only), DS5020 (FibreChannel only), DS5100 or DS5300 (FibreChannel only). Some legacy systems such as DS4700 can also be supported. These operating systems platforms can share SAN connectivity as can any supported operating systems. Zoning and partitioning tools should be used to separate dedicated access to volumes on the storage systems.

Supported File Sharing Environments

The above hardware configurations will allow you to have Linux, Apple, and Windows client systems share *physical* access to the same storage systems. When shared access to Files is required, the same basic hardware connectivity table applies. However, a specific set of file sharing mechanisms and SAN topology are required.

The solutions described in this guide are based upon StorNext software from Quantum. Apple Xsan is a SAN files sharing suite sold by Apple. It was developed to be fully compatible with Quantum StorNext (though it offers a very limited sub-set of the optional software add-on products offered by Quantum). Apple supports only a homogenous Apple environment with Xsan MDCs, Xsan clients, and Apple or Promise Disks. But by replacing the Xsan MDCs with StorNext MDCs a much more robust and flexible solution is possible.

Meta Data Controller	MDC Operating System	Supported Clients	Supported Client Operating System	Supported Storage Systems
Apple Xsan MDC	Mac OSX	Xsan clients only	Mac OSX Only	Apple or Promise Disk Only
Quantum MDC	Linux (Windows Server future support)	Xsan and/or StorNext clients	Mac OSX, Windows Server and Clients (XP, W7, Vista), Linux	Many, including IBM DS3000, DS4000, and DS5000

Note: A customer seeking connection that has already deployed Apple Xsan MDC will need to transition to a Quantum StorNext MDC running either on Linux or Windows Server. However the Apple Xsan clients can remain.

SOLUTION COMPONENT DESCRIPTIONS

SAN Architecture criteria

- Limit number of paths for Xsan/StorNext clients to four
 - This is to limit both complexity and the amount of time required to initialize clients on the SAN
- Redundant Meshed FibreChannel SAN (cross-connected)
 - Some IBM documentation may seem to preclude this, but those documents are based on an old multipathing methodology
 - Limits cause for failover to Controller failure
- Specific Firmware, software, and Operating System Levels

Apple Clients

- Apple connection without StorNext/Xsan support, or
- May use StorNext for shared SAN access
 - Apple Xsan Client software licensed from Apple on the Mac clients
 - Note: For this solution, you cannot use Apple servers as the Meta Data Controllers (MDCs). The reason is - if you use the Xsan clients and Xsan Meta Data Controller (MDC), then your primary support for the shared file system comes from Apple and the hardware solution described in this guide are NOT supported. However, if you use the StorNext MDC from Quantum, then Quantum will support the solution including the ability to connect Xsan clients on the shared SAN to non-Apple SAN hardware. As you can see, this is a much more flexible solution.
 - Must use Linux or Windows StorNext Meta Data Controllers licensed from Quantum
 - Must use supported ATTO HBA and MultiPath Director driver purchased from ATTO or ATTO Resellers

Windows Clients

- Windows client (XP, Vista, Windows 7) connection without StorNext Client software or
- May use StorNext software for shared SAN access
 - Must use Linux or Windows StorNext Meta Data Controllers with StorNext licensed from Quantum
 - Must use supported ATTO HBA and MultiPath Director driver purchased from ATTO or ATTO Resellers

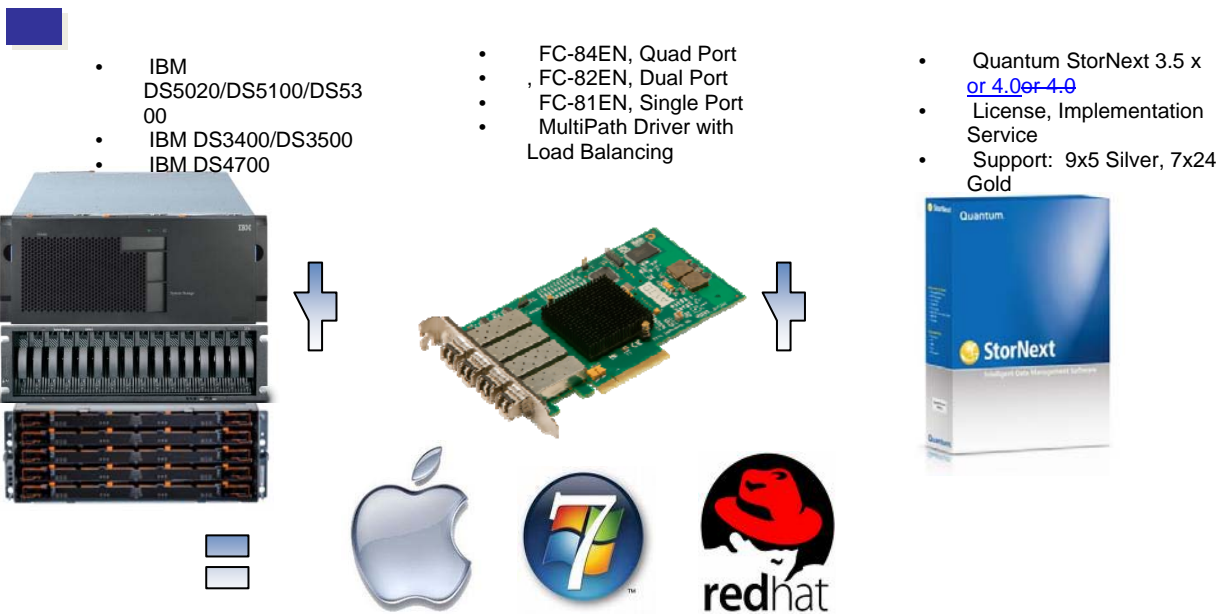
Windows Server MDCs (Future Support)

- Must buy Quantum StorNext software
- Must buy StorNext Installation services from Quantum
- Must buy StorNext Gold or Silver support contract from Quantum

- Must use supported ATTO HBA and MultiPath Director driver purchased from ATTO or ATTO Resellers

Linux StorNext Clients and MDCs

- Must buy Quantum StorNext software
- Must buy StorNext Installation services from Quantum
- Must buy StorNext Gold or Silver support contract from Quantum
- If the environment includes MacOS and/or Windows Clients (XP, Vista, or Windows 7) Must use supported ATTO HBA and MultiPath Director driver purchased from ATTO or ATTO Resellers
- If the environment is homogenous Linux, refer to the list of currently supported HBA's and multipath drives available for the operating system. This will include QLogic and Emulex HBA's.



IBM DS Storage

- DS5100/DS5300
- DS5020
- DS4700 (Legacy)
- DS3400
- DS3500
- Future

DS5100, DS5300

- FC, iSCSI connectivity
- (480) SSD, FC, SED

DS5020

- FC, iSCSI connectivity
- (112) FC, SED (DS5020)

DS3500

- FC, SAS, iSCSI connectivity

DS3200, DS3300, DS3400

- FC, SAS, iSCSI connectivity
- (48) SAS, SATA drives



SMB / Entry-level	Workgroup	Department	Data Center	HPC
<ul style="list-style-type: none"> • SAS, iSCSI connectivity • DAS/shared DAS/IP SAN • Simple management • Snapshot replication 	<ul style="list-style-type: none"> • iSCSI, FC connectivity • Shared DAS, small SAN • Simple management • Snapshot replication 	<ul style="list-style-type: none"> • FC, iSCSI connectivity • Homogeneous SANs • Performance value • Configuration flexibility • Local/remote replication 	<ul style="list-style-type: none"> • FC connectivity • Heterogeneous SANs • Highest performance • Configuration flexibility • Local/remote replication 	<ul style="list-style-type: none"> • FC connectivity • System clusters • GPFS • Maximum bandwidth • Capacity density

Burst I/O rate cache reads (IOPS)	110,000	140,000	174,000	200,000	650,000	700,000
---	---------	---------	---------	---------	---------	---------

Sustained I/O rate disk reads	21,000	30,000	35,000	50,000	76,000	170,000
-------------------------------------	--------	--------	--------	--------	--------	---------

Sustained I/O rate disk writes	4,500	7,500	11,000	9,500	20,000	38,000
--------------------------------------	-------	-------	--------	-------	--------	--------

Drives	48 SAS-1 /SATA	96 SAS-2 / NL- SAS / SED	96 SAS-2 / NL-SAS / SED	112 FC/SATA/ FDE	448 FC/SATA/ FDE	448 FC/SATA/FDE /SSD* (480) SATA*
--------	----------------------	--------------------------------	----------------------------------	------------------------	------------------------	--

	DS3400	DS3500	DS3500 Turbo	DS5020	DS5100 base performance	DS5300
Burst throughput cache read (MB/s)	1,600	2,500	6,000	3,700	1,600	6,400

Sustained throughput disk read (MB/s)	925	2,000	4,000	1,600	1,600	6,400
Sustained throughput disk write (MB/s)	720	1,100	2,400	1,400	1,300	5,300
Host ports	4	8	8	4	8	16

DS3500 Supported Drives



15K



10K



7200

Capacity	2.5-in.	3.5-in.
73 GB		
146 GB		
300 GB	FDE	
450 GB		
500 GB		
600 GB		FDE
1 TB		
2 TB		

Solutions

<p>Apple Hardware Connect Only</p> <p>Apple Servers and/or clients with no shared file access</p>	<p>Multiple Apple Host Server or Client connected to DS3000/DS4000/DS5000 Storage</p> <ul style="list-style-type: none"> Apple Leopard and Snow Leopard <ul style="list-style-type: none"> Using ATTO HBAs and ATTO MP Director Without StorNext or Xsan support This is basic host connectivity support.
<p>Apple Homogenous Shared File</p> <p>Apple Xsan Clients with Linux MDC (Meta Data Controller)</p>	<p>Multiple Apple Host Client connected to DS3000/DS4000/DS5000 Storage</p> <ul style="list-style-type: none"> Apple Leopard and Snow Leopard using Xsan Client software <ul style="list-style-type: none"> Using ATTO HBAs and ATTO MP Director With StorNext MDCs running on RHEL/SLES or Windows Server <ul style="list-style-type: none"> Using ATTO HBAs and ATTO MP Director
<p>Windows Client</p>	<p>Multiple Host Server or Client connected to DS3000/DS4000/DS5000 Storage</p>

Hardware Connect Only Windows clients with no shared file access	<ul style="list-style-type: none"> • Windows XP, Vista, Windows 7 <ul style="list-style-type: none"> • Using ATTO HBAs and ATTO MP Director • <u>Without StorNext or Xsan support</u> • This is basic host connectivity support.
Linux Homogenous Linux StorNext Clients with Linux MDC	<p>Multiple Linux Host Client connected to DS3000/DS4000/DS5000 Storage (NOW)</p> <ul style="list-style-type: none"> • RHEL/SLES using StorNext Client software <ul style="list-style-type: none"> • Using any supported HBA: QLogics, Emulex, or ATTO HBAs and ATTO MP Director • With StorNext MDCs running on RHEL/SLES <ul style="list-style-type: none"> • Using supported HBA: QLogics, Emulex, or ATTO HBAs and ATTO MP Director
Heterogeneous Apple Xsan Clients, Linux StorNext Clients, Windows Clients and Linux MDC	<p>Mixed Host Apple, Windows Client and Linux Client connected to DS3000/DS4000/DS5000 Storage</p> <ul style="list-style-type: none"> • Apple Leopard and Snow Leopard using Xsan Client software <ul style="list-style-type: none"> • Using ATTO HBAs and MultiPath Director • Windows XP, Vista, and Windows7 clients using Quantum StorNext Client software <ul style="list-style-type: none"> • Using ATTO HBAs and MultiPath Director • RHEL/SLES systems using StorNext Client software <ul style="list-style-type: none"> • Using ATTO HBAs and ATTO MultiPath Director • With StorNext MDCs running on RHEL/SLES <ul style="list-style-type: none"> • Using ATTO HBAs and ATTO MultiPath Director

APPLE, WINDOWS, and STORNEXT SUPPORT MATRIX

	O/S	Architecture	HBA	MPF	Supported
Meta Data Controller					
Linux	RHEL 5.3	x64	If clients are also Linux: supported HBA: QLogics, Emulex, If clients are MacOS, Windows Client, or mixed: Supported <i>ATTO HBAs</i>	RDAC, MPP, or ATTO as appropriate	Yes
Linux	RHEL 5.3	x64	If clients are also Linux: supported HBA: QLogics, Emulex, If clients are MacOS, Windows	RDAC, MPP, or ATTO as appropriate	

			Client, or mixed: Supported <i>ATTO HBAs</i>		
Linux	SLES 10.2	x64	If clients are also Linux: supported HBA: QLogics, Emulex, If clients are MacOS, Windows Client, or mixed: Supported <i>ATTO HBAs</i>	RDAC, MPP, or ATTO as appropriate	Yes
Linux	RHEL 5.3	x64	If clients are also Linux: supported HBA: QLogics, Emulex, If clients are MacOS, Windows Client, or mixed: Supported <i>ATTO HBAs</i>	RDAC, MPP, or ATTO as appropriate	
Client					
Apple OSX	Leopard	x64	<i>Supported ATTO HBAs</i>	ATTO	Yes
Apple OSX	Snow Leopard	x64	<i>Supported ATTO HBAs</i>	ATTO	Yes
Linux	RHEL 5.3	x86, x64	If clients and MDC are all Linux: supported HBA: QLogics, Emulex, If mixed with clients that are MacOS, Windows Client: Supported <i>ATTO HBAs</i>	ATTO	Yes
Linux	SLES 10.2	If clients and MDC are all Linux: supported HBA: QLogics	Emulex	ATTO	Yes
Windows	2003 SP2	x86, x64	<i>Supported ATTO HBAs(future)</i>	MPIO/DSM, ATTO	
Windows	2008 SP2	x86, x64	<i>Supported ATTO HBAs(future)</i>	MPIO/DSM, ATTO	

Windows	2008 R2	x86, x64	Supported ATTO HBAs(future)	MPIO/DSM, ATTO)	
Windows	Windows 7	x86	Supported ATTO HBAs	ATTO	
Windows	Vista	x86	Supported ATTO HBAs	ATTO	
Windows	XP	x86	Supported ATTO HBAs	ATTO	
Mixed			As appropriate above	As appropriate above	

SOLUTION RESTRICTIONS AND RECOMMENDATIONS

There are several important aspects to the configuring of the SAN that must be understood. The first is that the SAN design is architected to avoid logical drive failover in the storage subsystem. This is accomplished by employing a meshed fabric where each HBA on the server can see both controllers on the storage subsystem. The second aspect has to deal with access to the shared logical drives. There are three types of data to store on the shared logical drives; Metadata, Journal, and Data. All systems (MDCs and Clients) must be able to see the Data LUNs. The MDCs must also be able to both see the Metadata and Journal LUNs. It is important that the clients do not have access to the Metadata or Journal LUNs. When configuring the storage partitioning within the storage subsystem, the methodology is as follows:

- All HBAs in the MDCs will be defined within a single host definition (called MDCs?). This host definition will have the Metadata and Journal LUNs directly mapped to this host definition.
- All Clients have their own host definitions. If desired, all clients that have the same OS can have their HBAs combined within a single host definition.
- All Clients and the MDCs will be assigned to a single host group (called StorNext?). All of the Data LUNs will be assigned to that host group.

	HBAs	Connections	Zone
MDC	Single or multiple (typical)	Up to four	All ports see Controller A and Controller B on DSxxxx
<p>MDC - Limit number of defined paths to optimize initialization time. There are few MDCs in comparison to clients, so this choice should not be as restrictive as below. This is not a hard requirement, but is recommended.</p> <p>Most configurations consist of two MDCs for high availability. Both MDCs must have the same configuration; OS, HBAs, Driver Levels, etc</p>			
Client	Single (typical) or multiple	Up to two	All ports see Controller A and Controller B on DSxxxx
<p>Client – Limit number of defined paths to four to optimize initialization time. This is not a hard requirement, but is recommended.</p>			

Storage	Controller A	Up to 8	All ports see all MDC and client ports
	Controller B	Up to 8	All ports see all MDC and client ports

RECOMMENDED RAID CONFIGURATIONS

	RAID level	RAID Array size	Segment size	Disk type
Journal	1 or 10	As required	256	FC or SAS recommended
Meta Data	1 or 10	As required	256	FC or SAS recommended
<ul style="list-style-type: none"> Journal and Meta Data may reside in separate LUNS carved from the same RAID Array if required in small implementations. However, the ideal condition is separate LUNS tied to separate RAID groups for Journal and Meta Data. If possible, these two LUNS should be on a different storage system than the data LUNS. StorNext can stripe across multiple LUNS. It is recommended that multiple RAID 1 Groups be used for Journal and Meta Data rather than similarly sized RAID 10s. This allows multiple threads to be talking to disk offering enhanced performance. 				
Data	5	4+1	256	As required to meet performance criteria
	5	8+1	128	As required to meet performance criteria
<ul style="list-style-type: none"> Data – Based upon StorNext file system header block size of 1MB, it is important to keep this in mind when selecting the segment size. RAID Array (also called Group) size and appropriate segment size should result in stripe size of 1MB. This will result in optimizing write throughput speeds. Ex. RAID5 4+1 (4 data) drives x 256K (seg size) = 1 MB Write stripe & RAID5 8+1x128K (seg size) = 1 MB stripe. Note – this rule does not apply to all deployments and could be dramatically different for oil and gas industry deployments among others. It is very important to understand the I/O block size of the applications sitting atop the StorNext File system. 				
Data	6	4+2	256	As required to meet performance criteria
	6	8+2	128	As required to meet performance criteria
<ul style="list-style-type: none"> Some customer applications may specify SATA drives and/or require additional RAID protection. The DS3000/DS4700/DS5000 storage servers offer additional levels of RAID protection if desired. 				

SOLUTION INSTALLATION

1. Use a meshed fabric with each HBA zoned to see both controllers but controllers unable to see each other.
2. Use RDAC or MPIO failover drivers as appropriate for your operating systems (ATTO for Apple and/or Windows XP, Vista, or Windows 7). Do not use the StorNext failover driver
3. For Installations that include Apple (MacOS) or Windows (XP, Vista, Windows 7) Clients
 - a. Set host type **MacOS**
 - b. This will :
 1. Set **TPGS** On
 2. Set **AVT** Off
 - c. MAC's and Windows XP, Windows Vista, and Windows 7, must use ATTO Celerity FC81-EN, FC82-EN, or FC-84-EN HBAs

Atto site
<http://www.attotech.com/>

Driver downloads,
<https://www.attotech.com/register/index.php?FROM=http://www.attotech.com/downloads.html>

The user is required to register and create a free account. Then navigate to,
Fibre Channel HBAs for Storage Partners.
IBM DS Series Celerity FC-8xEN
4. When connecting to MAC systems
 - a. Switched environments preferred
 - b. In non-switched (direct-connect) environments, configure the ATTO host adapter for an "Arbitrated Loop" connection instead of allowing them to auto negotiate.
 - c. Use Host Region "Mac OS."
5. For Linux Clients mixed with Apple Clients and Windows Clients
 - a. If Apple servers or clients are sharing physical connection to a storage system, it is important to use the proper host type definitions.
 - b. Use Host Region "Mac OS."
6. How to handle mixing multiple server and client O/S types when they require access same LUNS:
 - a. The MDCs in the approved solutions will always be Linux or Windows
 - b. Clients can be Apple, Windows or Linux
 - c. To allow disparate O/S host connections:
 1. Create a Host Group with no Host type associated
 2. Use Host Region "Mac OS."
7. In some StorNext environments (DS5020 and DS4700) it is recommended that you do not have a LUN 0 mapped. There may be confusion with StorNext thinking that LUNs off of the controller with LUN 0 mapped will not be compatible because it thinks AVT is off and it is needed ON.
8. Recommended cache settings: Set pre-fetch value to any number other than 0 (zero) to enable automatic prefetch.

9. For clients, limit the number of defined paths to four. This gives you the redundancy you need while minimizing boot time. During boot, StorNext tries EVERY defined path.

10. RAID 1 or 10 for Journal and Meta Data
 - a. Multiple RAID 1 Groups would be preferable to a similarly size RAID 10
 - b. Placing Meta Data and Journal on a storage system separate from the data is preferable.
11. It is desirable to separate Meta data and Journal LUNS from the Data LUNS.
 - a. In small installations, Meta data, Journal and Data can be different LUNS on the same RAID group
 - b. However, it is best to separate Meta data and Journal into their own RAID 1 RAID group.
 - c. When Data gets to be large (approximately 250,000 files or more), it is best to completely separate Meta Data and Journal into their own storage system.
 - d. A typical segment size for the Metadata and Journal (Segment Size is assigned on a LUN-by-LUN basis) is 256. Though this may vary with various applications. It is critical to understand the I/O demands of the applications being deployed.
12. The size and RAID type used for Data are really dependent upon customer requirements. But these levels are recommended for typical media applications:
 - a. RAID 5 4+1 or RAID 5 8+1
 - b. Explanation: RAID 5 4+1 means a RAID group using RAID 5 with 4 data drives and one parity drive. Even though in actuality, the data and parity would be striped across all drives.
13. Some applications (especially those using SATA drives) use different RAID types. Many scientific applications might keep very large data stores on relatively inexpensive SATA disks. In these cases, RAID 6 may be desired.
 - a. RAID 6 4+2 or RAID 6 8+2
14. Typically, the media applications using StorNext write in 1M block. So:
 - a. Using the formula Block size / number of data disks
 - b. RAID 5 4+1 using 256K segment size
 - c. RAID 6 4+2 using 256K segment size
 - d. RAID 5 8+1 using 128K segment size
 - e. RAID 6 8+2 using 128K segment size
 - f. These segment sizes are representative of the 1 meg I/O size typical in most media industry applications. However, some applications (notably those in the oil and gas industry) may require different segment sizes. It is critical to understand the I/O demands of the applications being deployed.
15. Apple Clients HBAs: ATTO HBAs use default settings
16. Linux Clients HBAs: ATTO HBAs use default settings.
17. Linux MDC and Clients HBAs: ATTO HBAs use default settings.
18. Windows Clients HBAs: ATTO HBAs use default settings.
19. Use standard performance tuning documentation as regards to disk and RAID group layouts etc.

QUANTUM / APPLE JOINT PRODUCT SUPPORT PROGRAM



Quantum / Apple Joint Product Support Program

Overview

Quantum and Apple have agreed to a cooperative support relationship for Quantum's StorNext and Apple's Xsan products. This means joint customers are assured that both companies will work together to resolve any technical issues with our respective products.

Xsan Controlled SAN File System

In the case of the joint customer having an Xsan controlled SAN file System, (i.e. Xsan runs the main server or Meta Data Controller), then Apple should receive the first call. If Quantum does receive the first call Quantum will start a case report and try to determine if the issue is related to the QUANTUM portion of the install. If after some investigation it is determined that an Apple case remedy is required, then Quantum shall notify the Apple Technical Support Contact, and shall provide the customer's information. Apple shall contact the joint customer in accordance with that party's customer support and/or maintenance agreement with the joint customer. Quantum shall continue to be available to provide information and/or assistance, as reasonably necessary.

StorNext Controlled SAN File System

In the case of the joint customer having an StorNext controlled SAN file system (i.e. StorNext runs the main server or Meta Data Controller), then Quantum should receive the first call. If Apple does receive the first call Apple will start a case report and try to determine if the issue is related to the Apple portion of the install. If after some investigation it is determined that an Quantum case remedy is required, then Apple shall notify the Quantum Technical Support Contact, and shall provide the customer's information. Quantum shall contact the joint customer in accordance with that party's customer support and/or maintenance agreement with the joint customer. Apple shall continue to be available to provide information and/or assistance, as reasonably necessary.

Cooperation Effort Between Apple and Quantum

If, after making reasonable but unsuccessful efforts to provide a case remedy, the parties shall then endeavor to share information and use cooperative efforts to determine and/or isolate the cause of the case report. The party in whose product the cause is isolated shall deliver or implement a timely case remedy; the other party agrees to use reasonable efforts to supply information, assist, and confirm the case remedy. Apple and Quantum are responsible for support and maintenance of its own products, and are not authorized to support or maintain the products of the other party.

Closing a Case

The consent of a joint customer shall be required to close a case.

MIGRATING FROM XSAN MDC TO STORNEXT MDC

Q. Are there new features in Xsan 2.2 that make it harder to swap out XServe MDCs and replace them with StorNext Linux MDCs?

A: "NamedStreams" is the Apple Xsan feature that causes some problems with StorNext. If you don't turn it on, you're fine. By default, "NamedStreams" is turned on. NamedStreams will make things like the basic StorNext command line copy command, cvcp, not work. Before Xsan 2.2, you could have it copy the "." files, but now you can't, so you're stuck doing a move with the Apple Finder (which is basically as good as cvcp anyway). Other than that, it shouldn't be much of an issue since it doesn't affect clients that don't know about it, and by definition Windows/Linux clients don't know about resource forks.

Q. Does a customer need to buy A special level of StorNext support to support StorNext in Xsan environments? Is this true whether you are using SN MDCs or Apple MDCs?

A: No, You don't need to buy any special support to get support for StorNext in an Xsan environment. See next page.

ATTO Technology Inc.



ATTO Technology, Inc. Value Proposition

ATTO Technology is a leader in storage and connectivity solutions. In business since 1988, ATTO has maintained a consistent strategy of technical innovation by continually creating new and exciting Host Bus Adapter (HBA) products that are on the leading edge of the storage market. Throughout the years, ATTO has maintained a reputation for providing high-performance products in a variety of storage connectivity technologies for a wide variety of market segments.



ATTO Celerity Fibre Channel HBAs

- **Advanced Architecture** — Don't settle for just any host bus adapter. Celerity Fibre Channel host bus adapters incorporate industry-standard components into a unique architecture to deliver a robust and scalable connectivity solution. Celerity offers a highly integrated family of adapters designed to transfer data efficiently and flawlessly using a combination of features that enhance the data delivery process. ATTO's attention to detail in board design, specifically signal quality, ensures reliable connectivity. This exceptional design minimizes data corruption, reduces transmission errors, and protects other system components from harmful interference.

Celerity Fibre Channel HBAs help solve advanced storage connectivity challenges by offering high-end features such as multipathing for failover and load balancing in clustered applications with select storage vendors.

- **Industry Proven Performance** — ATTO's commitment to Fibre Channel technology and storage area networking began over 15 years ago and continues with first-to-market products that include the latest performance levels and features. ATTO Celerity Fibre Channel HBAs deliver performance, reliability and connectivity with a wide range of port options, OS support and proven industry certifications and interoperability.

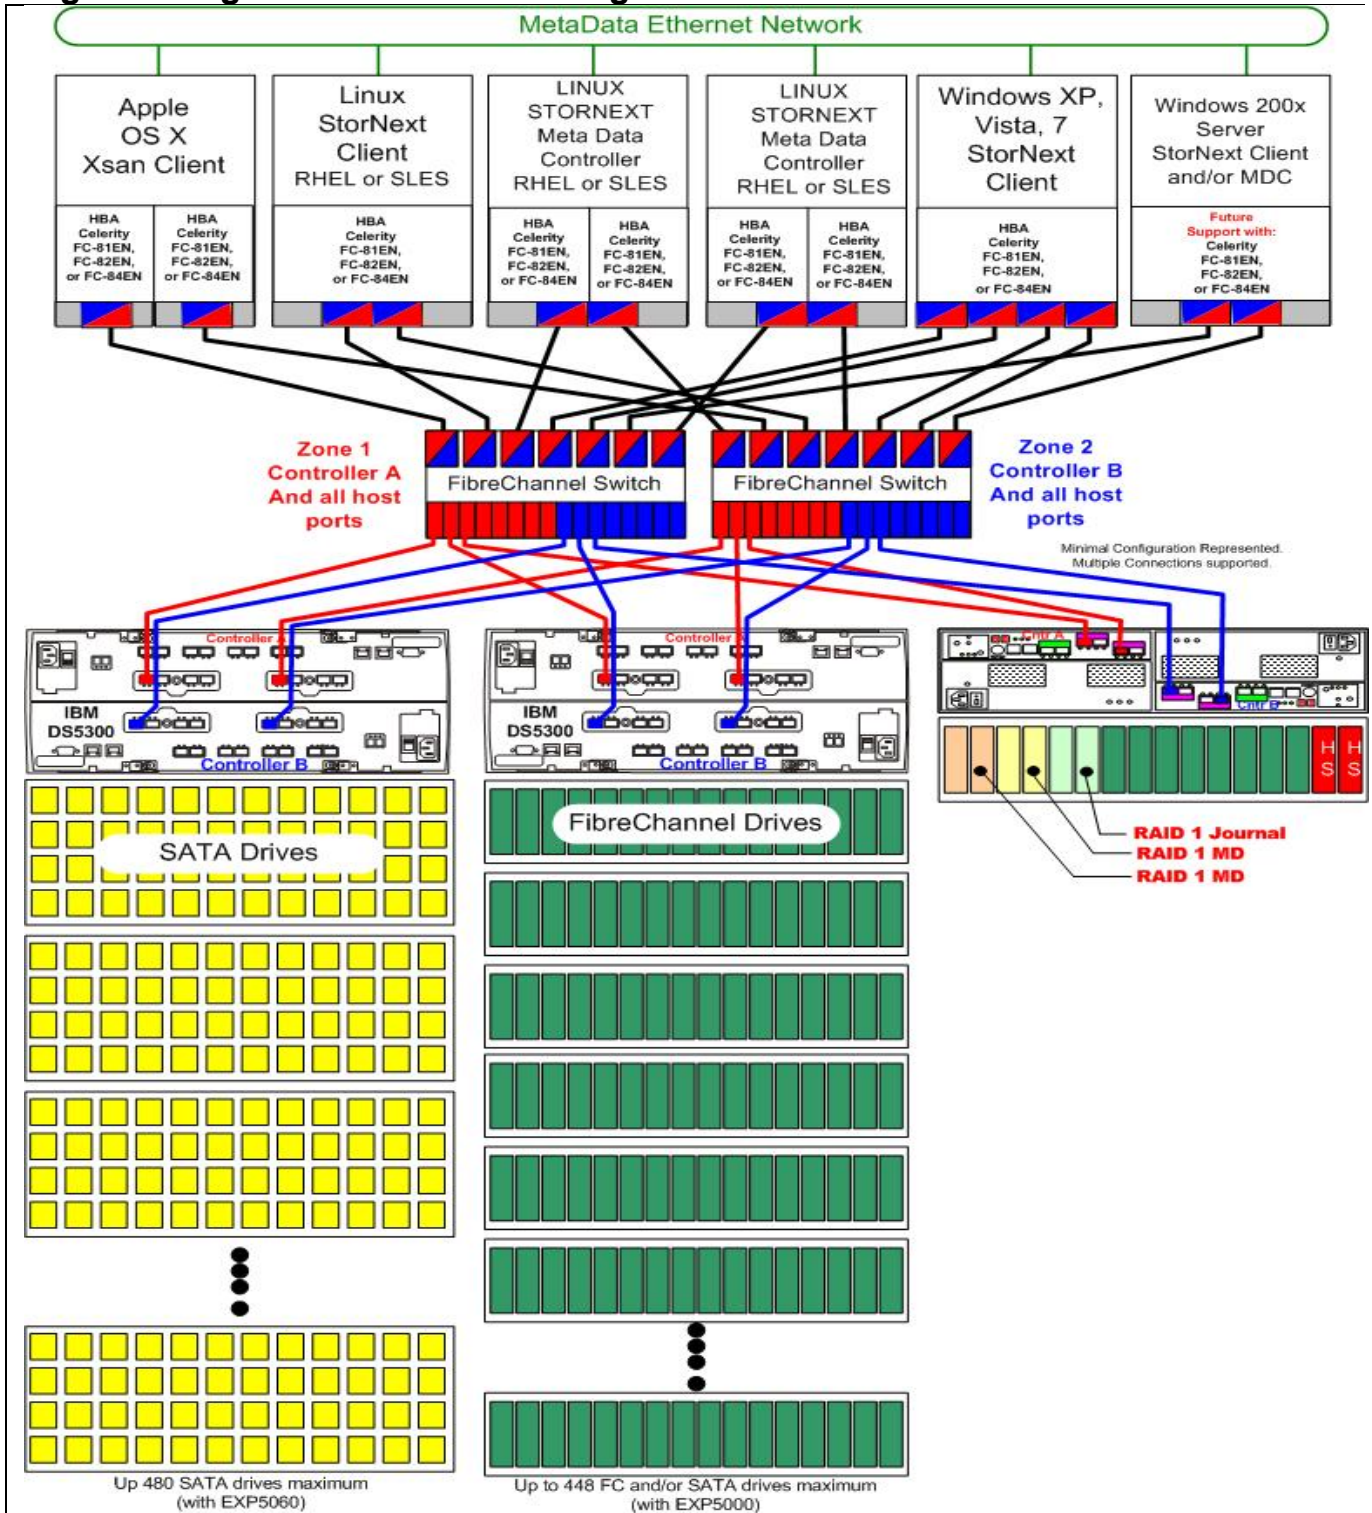
ATTO Celerity 8Gb Fibre Channel HBAs

Product Features	Celerity HBAs		
	FC84EN	FC82EN	FC81EN
Fibre Channel Ports	4	2	1
Maximum Data Rate	8Gb	8Gb	8Gb
Maximum Transfer Rate (Full Duplex)	6.4 GB/s	3.2 GB/s	1.6 GB/s
Maximum Transfer Rate (Half Duplex)	3.2 GB/s	1.6 GB/s	800 MB/s
Bus Type	PCIe 2.0	PCIe 2.0	PCIe 2.0
Bus Characteristics	X8	X8	X8
Optical Interface	SPF+LC	SPF+LC	SPF+LC
Maximum Cable Length	300m-2GB 150m-4GB 50m-8GB	300m-2GB 150m-4GB 50m-8GB	300m-2GB 150m-4GB 50m-8GB
Low Profile Form Factor	NO	YES	YES
Advanced Data Streaming (ADS)	YES	YES	YES
Software RAID Support	YES	YES	YES
Developers Kit (Target Mode and API)	YES	YES	YES
Windows (Server)	YES	YES	YES
Windows (work station / client)	YES	YES	YES
Linux (Red Hat, SUSE)	YES	YES	YES
MAC OSX	YES	YES	YES
VMWare ESX	YES	YES	YES
RoHS Compliant	YES	YES	YES

http://www.attotech.com/selectioncharts/Celerity_Selection_Chart.pdf

APPENDIX A: EXAMPLE CONFIGURATIONS

Large Heterogeneous StorNext configuration



Example 1. Drawing depicts large and expandable SAN with heterogeneous client connectivity. Multiple DS5300 connected to meshed SAN fabric with a minimum of two Fibre Channel switches. For DS5100/DS5300 configurations, a minimum of two switches is recommended. Drawing shows multiple DS5300 for Data storage and a separate DS5200 for Meta Data.

Detailed explanation of example 1.

Example Large StorNext configuration showing DS5300's with Fibre Channel and/or SATA disk for Data and a separate DS5020 for Meta Data and Journal. These examples are representative only and should not be seen as restrictive. The idea is to use our knowledge of these products to craft solutions with the proper performance and resiliency characteristics to meet customer needs and budget concerns.

DS5020 for Journal and Meta Data with FC drives

- This is an example configuration only
- RAID 1 for StorNext Meta Data
- RAID 1 for StorNext Journal
- Add additional RAID 1's rather than growing to RAID 10 for performance reasons.
- Controller A can see all MDC host ports but not Controller B
- Controller B can see all MDC host ports but not Controller A

DS5300's with FC and/or SATA for Data

- This is an example configuration only
- Multiple RAID 5 (both 4+1 and 8+1) RAID Arrays for StorNext data
- Multiple RAID 6 (both 4+2 and 8+2) RAID Arrays for StorNext data may also be desirable in some situations
- Controller A can see all client and MDC host ports but not Controller B
- Controller B can see all client and MDC host ports but not Controller A

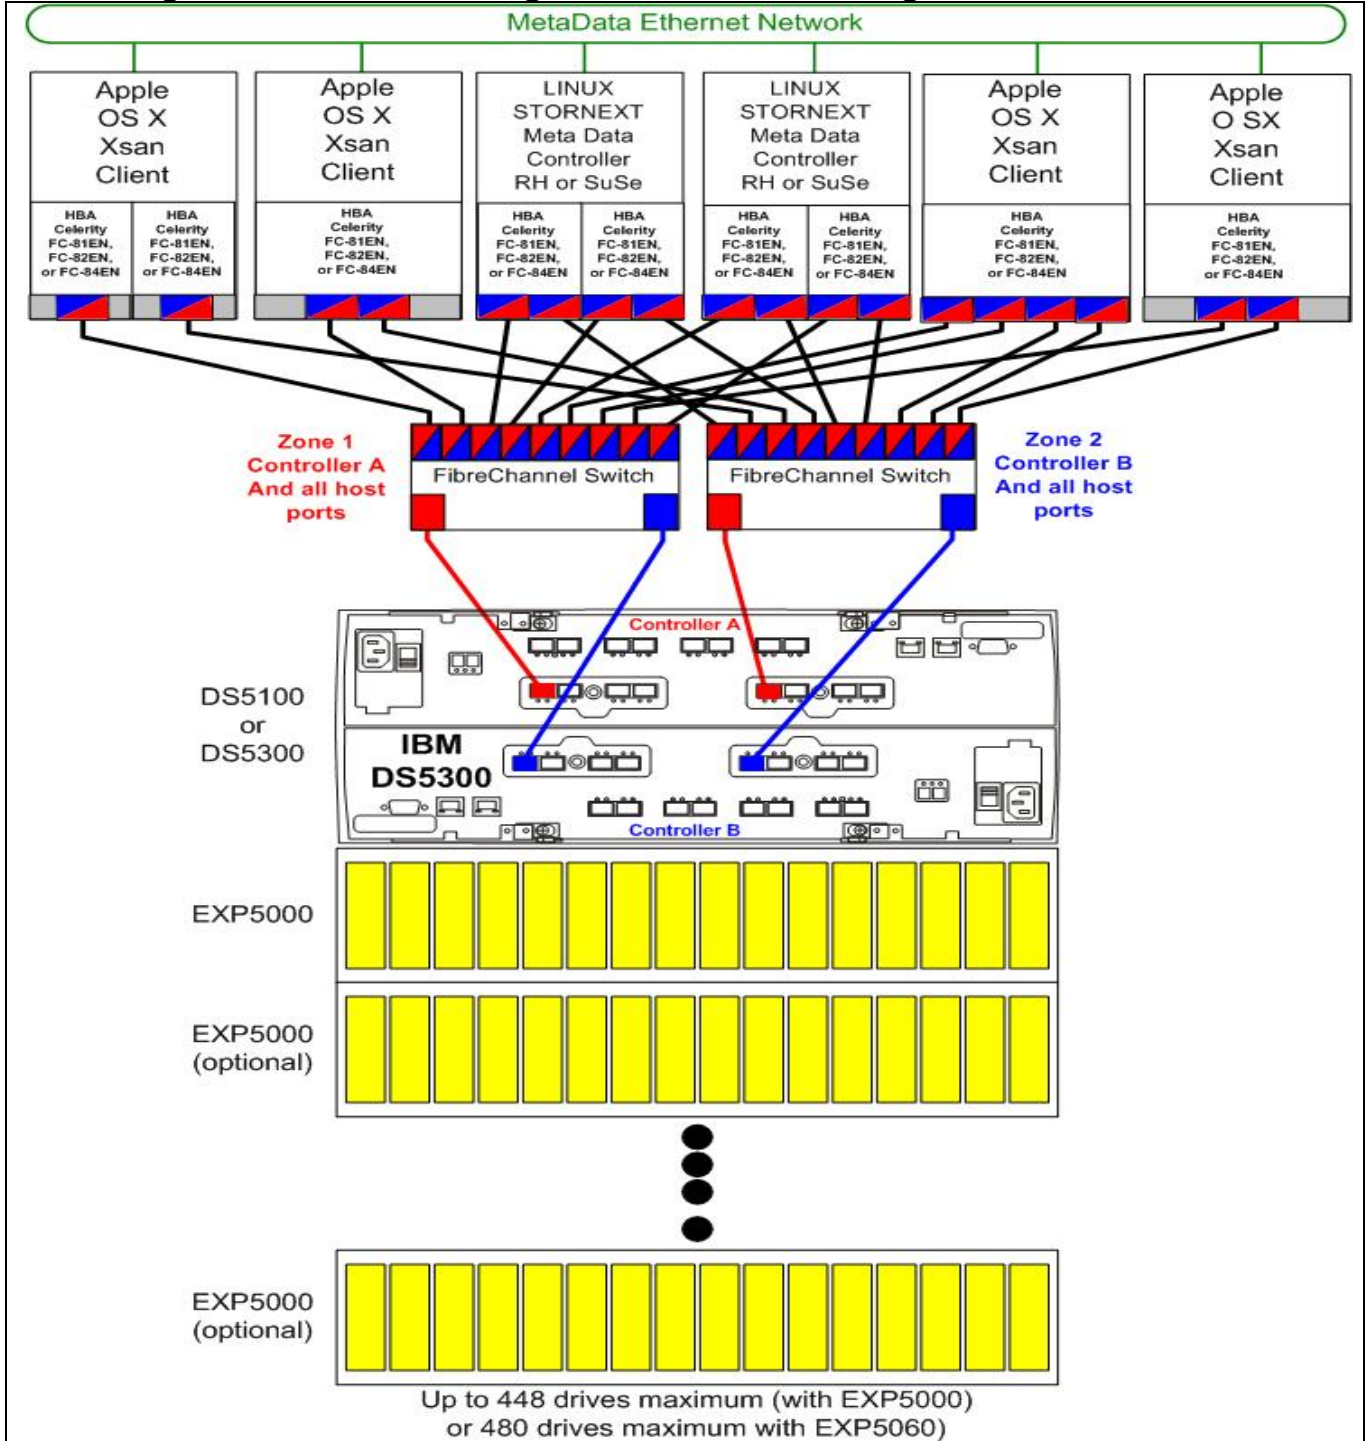
Mixed Clients – SAN Zoning

- Clients have single dual ported HBAs with connections into a redundant switched network and zoned to be able to see both Controller A and B
- MDCs each have two dual ported HBAs with connections into a redundant switched network and zoned to be able to see both Controller A and B

All Storage Configurations for StorNext environments

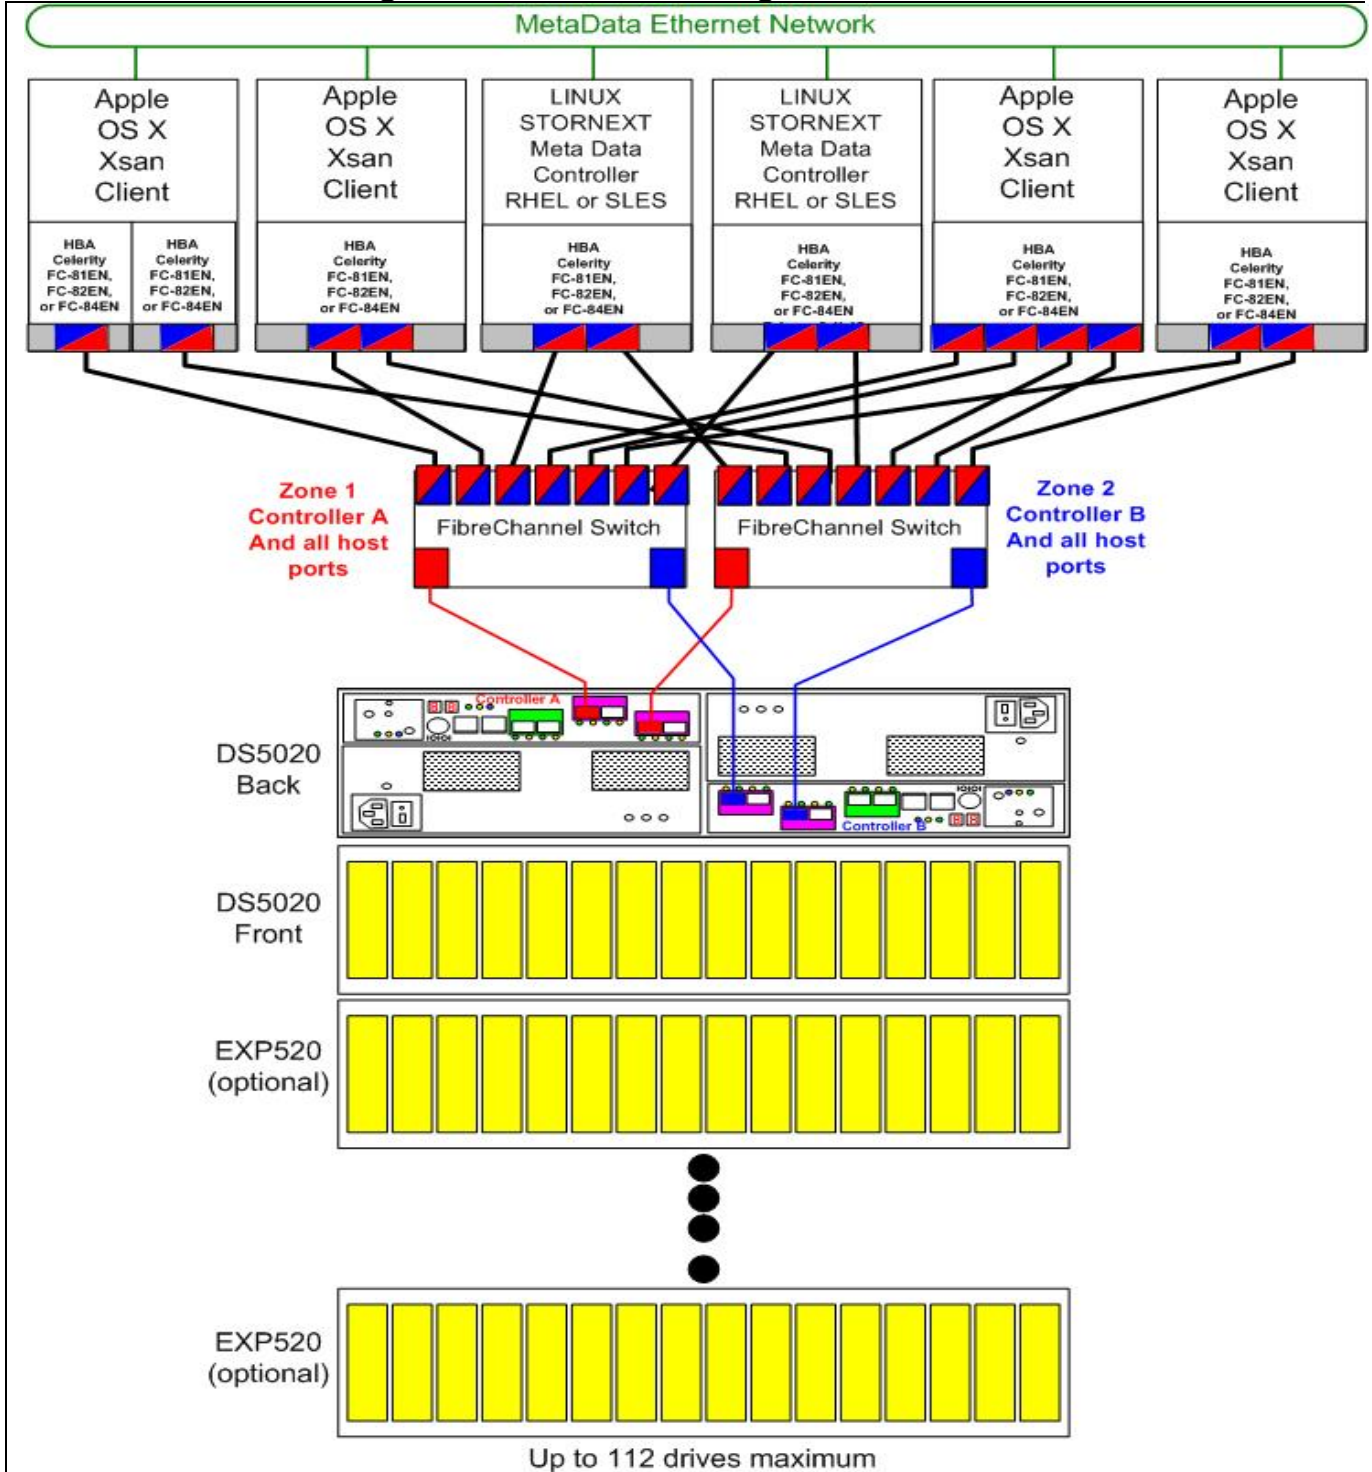
- AVT should be set OFF for all partitions and host types
- For **MAC Clients** there is now a host type definition "MAC OS" for the Clients. This was not previously available, and because of this a work-around had been developed for use until the MAC host type was available.
 1. Set **host type to MacOS, the previous workaround was to set it to LNXCLVMWARE**
 2. With the workaround, you also had to manually set **TPGS on** (via a provided script)
- For **Linux Clients (and VMware servers)**
 1. Linux Clients should use host type LNX. VMware Clients should use host type VMware,
 2. Previous versions of FW required the MAC hosts to use LNXCLVMWARE as the host type. This could cause conflicts if Linux and VMware servers also existed in the environment, which is why more host types have been added. Rules for the old workaround include:
If no MAC clients exist, then the host type LNX should be used, as this host type has AVT disabled.
If MAC hosts exist in the environment, Linux clients (and VMware servers) should use the host type LINUX,
 3. Insure that AVT is turned off for all host regions, including LINUX (via a provided script). This can be done by checking the bottom of the complete (all tab) system Profile.
- For mixed environments:
 1. Create a Host Group for all MDCs and Clients that will access the StorNext file system. There is no Host type associated with Host Groups (only individual hosts)
 2. Assign multiple host types to this group.

Mid to large -sized “MAC Homogenous” StorNext configuration



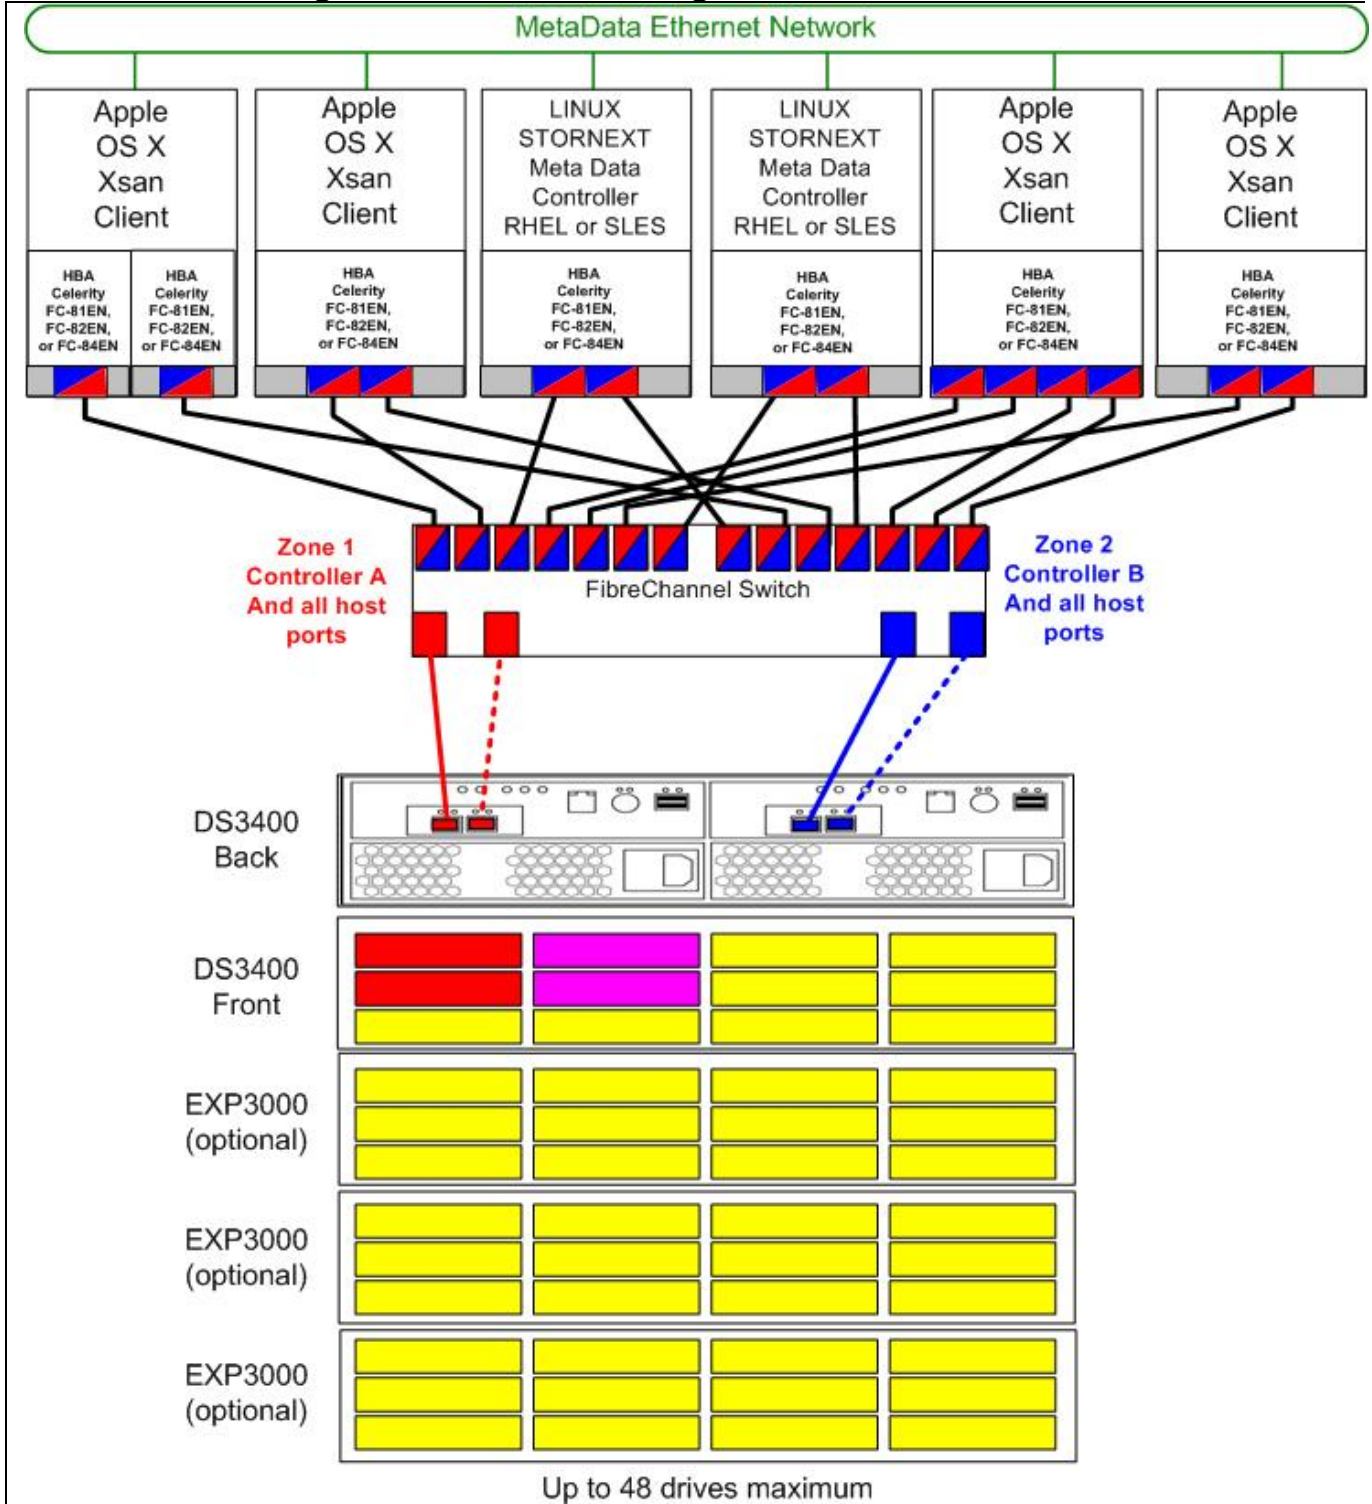
Example 2. In this example diagram Linux MDC Servers are shown, but only MAC Xsan Clients are shown. DS5300 connected to meshed SAN fabric with a minimum of two Fibre Channel switches. For DS5100/DS5300 configurations, a minimum of two switches is recommended. Drawing shows minimal SAN to DS5300 connections to simplify the drawing. However, up to 16 connections is supported. In this example, the Meta data Controllers each have a two dual ported HBA. Multiple multi-ported HBAs is not required, but it is recommended. Also in this example multiple HIC cards on the DS5300 are used to increase resiliency.

Mid-sized “MAC Homogenous” StorNext configuration



Example 4. In this example diagram Linux MDC Servers are shown, but only MAC Xsan Clients are shown. DS5020 connected to meshed SAN fabric with a single Fibre Channel switch. For DS5020 configurations, Multiple Fibre Channel switches are not required, but they are recommended. Drawing shows minimal SAN to DS5020 connections to simplify the drawing. However, up to 8 connections are supported. In this example, the Meta data Controllers each have a single dual ported HBA. Also in this example multiple HIC cards on the DS5020 are used to increase resiliency.

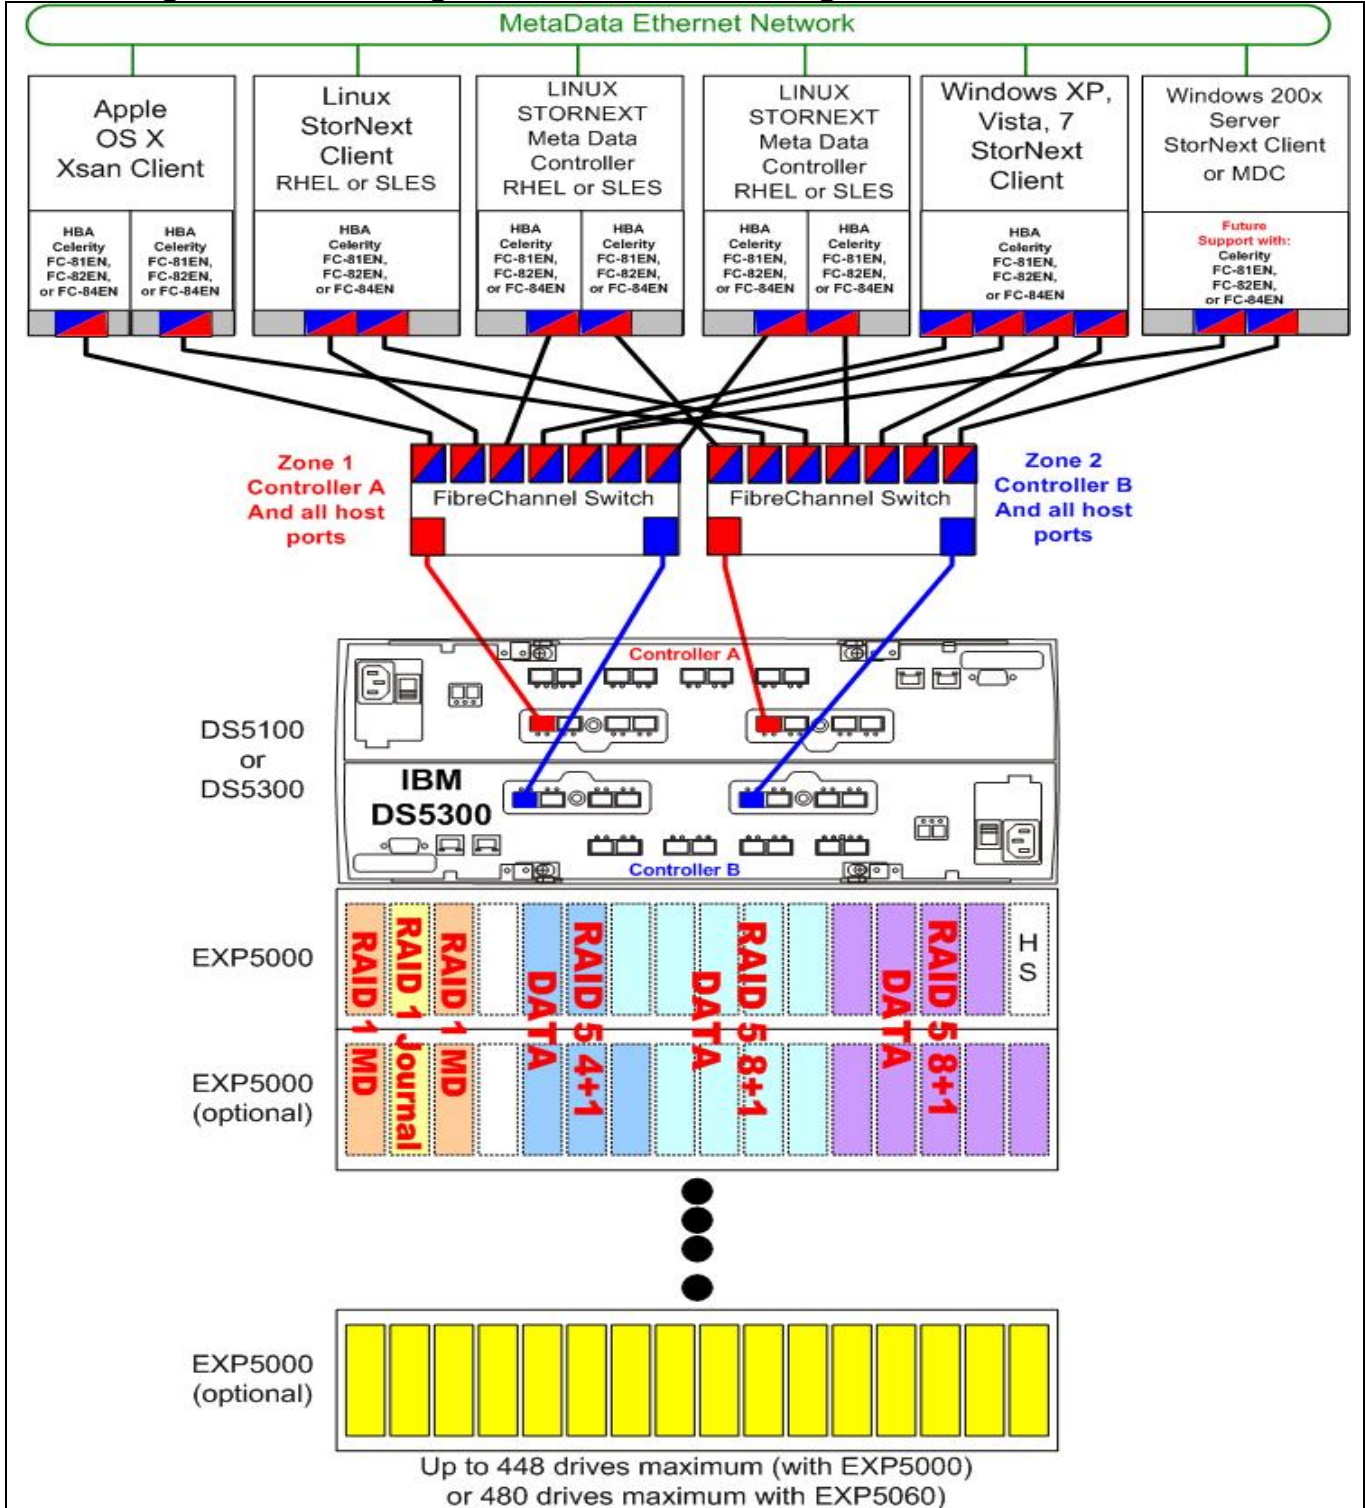
Small "MAC Homogenous" StorNext configuration



Example 5. In this example diagram Linux MDC Servers are shown, but only MAC Xsan Clients are shown. DS3400 connected to meshed SAN fabric with a single Fibre Channel switch. For DS3400 configurations, Multiple Fibre Channel switches are not required, but they are recommended. Drawing shows minimal SAN to DS3400 connections to simplify the drawing. However, up to 4 connections are supported. In this example, the Meta data Controllers each have a single dual ported HBA.

RAID 1 pairs for Meta Data and Journal are represented. Global Hot spares would be assigned as appropriate. While SAS drives would be recommended for Meta Data and Journal, SAS and/or SATA drives could be used for data as appropriate.

Mid to large -sized Heterogeneous StorNext configuration



Example 6. Drawing depicts mid to large size and expandable SAN with heterogeneous client connectivity. DS5300 connected to meshed SAN fabric with a minimum of two Fibre Channel switches. For DS5100/DS5300 configurations, a minimum of two switches is recommended. Drawing shows minimal SAN to DS5300 connections to simplify the drawing. However, up to 16 connections is supported.

In this example, the Meta data Controllers each have a single dual ported HBA. Also in this example multiple HIC cards on the DS5300 are used to increase resiliency. This drawing includes RAID Group type and size suggestions. However performance tuning guidelines should be used to plan RAID group layout between trays to maximize performance.

APPENDIX B: ATTO HBA AND MULTIPATH DIRECTOR INSTALLATION AND TROUBLESHOOTING MANUAL

- Please refer to ATTO Celerity MultiPath Director Manual for installation and troubleshooting.
 - http://www.attotech.com/software/files/manuals/Manual_Celerity_MultiPath-Director.pdf

Additional Reference Links

- IBM DS Series Portal: www.ibmndsseries.com
- ATTO Technology, Inc: www.attotech.com/solutions/ibm
- Quantum StorNext: www.stornext.com
- Apple Xsan: <http://www.apple.com/Xsan/>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information about the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of other companies:

Apple, the Apple logo, Mac, Mac OS, and Macintosh, are trademarks of Apple in the United States, other countries, or both.

Quantum, the Quantum logo and are registered trademarks of Quantum Corporation and its affiliates in the United States, other countries, or both.

ATTO, the ATTO logo are trademarks of ATTO Technology Inc, in the United States, other countries, or both.

Microsoft, Windows Server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Free Manuals Download Website

<http://myh66.com>

<http://usermanuals.us>

<http://www.somanuals.com>

<http://www.4manuals.cc>

<http://www.manual-lib.com>

<http://www.404manual.com>

<http://www.luxmanual.com>

<http://aubethermostatmanual.com>

Golf course search by state

<http://golfingnear.com>

Email search by domain

<http://emailbydomain.com>

Auto manuals search

<http://auto.somanuals.com>

TV manuals search

<http://tv.somanuals.com>