EMC® Host Connectivity Guide for Apple/Mac OS X

Contents

Figures

	Title	Page

EMC Host Connectivity Guide for Apple/Mac OS X

As part of an effort to improve and enhance the performance and capabilities of its product line, EMC from time to time releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all revisions of the software or hardware currently in use. For the most up-to-date information on product features, refer to your product release notes.

If a product does not function properly or does not function as described in this document, please contact your EMC representative.

This guide describes the features and setup procedures for Apple MAC OS X to EMC Symmetrix and CLARiiON storage arrays over Fibre Channel and iSCSI. The guide also lists the features and recommended configurations to deploy Xsan on the Apple MAC OS X platform.

Audience   This guide is intended for use by storage administrators, system programmers, or operators who are involved in acquiring, managing, or operating Symmetrix, CLARiiON, and host systems.

Readers of this guide are expected to be familiar with the following topics:

◆ Symmetrix or CLARiiON system operation

◆ Apple MAC OS X operating system with Xsan FS

Related documentation   For the most up-to-date information, always consult the *EMC Support Matrix* (ESM), available through E-Lab Interoperability Navigator (ELN) at http://elabnavigator.EMC.com, under the **PDFs and Guides** tab.

More information can be found in the following related documentation, available on Powerlink:

◆ EMC Connectrix B Series hardware reference manual*s*

◆ EMC CLARiiON documentation

◆ EMC Symmetrix documentation

Refer to www.attotech.com for more information on ATTO Host Bus Adapters or the ATTO Configuration Tool.

A web page dedicated to EMC customers, including information on various supported drivers and download information, is available at www.attotech.com/solutions/emc.

Refer to www.apple.com for more information on the Mac OS Leopard, Xsan, and other Apple products.

Conventions used in this guide

EMC uses the following conventions for notes, cautions, warnings, and danger notices.

Note: A note presents information that is important, but not hazard-related.

⚠ **CAUTION**

A caution contains information essential to avoid data loss or damage to the system or equipment. The caution may apply to hardware or software.

Typographical conventions

EMC uses the following type style conventions in this guide:

Normal	Used in running (nonprocedural) text for:
	• Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus)
	• Names of resources, attributes, pools, Boolean expressions, buttons, DQL statements, keywords, clauses, environment variables, filenames, functions, utilities
	• URLs, pathnames, filenames, directory names, computer names, links, groups, service keys, file systems, notifications
bold	• User actions (what the user clicks, presses, or selects)
	• Interface elements (button names, dialog box names)
	• Names of keys, commands, programs, scripts, applications, utilities, processes, notifications, system calls, services, applications, and utilities in text

italic	• Book titles • New terms in text • Emphasis in text
Courier	• Prompts • System output • Filenames • Pathnames • URLs • Syntax when shown in command line or other examples
Courier, bold	• User entry • Options in command-line syntax
Courier italic	• Arguments in examples of command-line syntax • Variables in examples of screen or file output • Variables in pathnames
<>	Angle brackets for parameter values (variables) supplied by user.
[]	Square brackets for optional values.
|	Vertical bar symbol for alternate selections. The bar means or.
...	Ellipsis for nonessential information omitted from the example.

Where to get help EMC support, product, and licensing information can be obtained as follows.

Product information — For documentation, release notes, software updates, or for information about EMC products, licensing, and service, go to the EMC Powerlink website (registration required) at:

http://Powerlink.EMC.com

Technical support — For technical support, go to EMC Customer Service on Powerlink. To open a service request through Powerlink, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or to answer any questions about your account.

Your comments Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send your opinion of this guide to:

techpub_comments@EMC.com

1

Supported Environments

This chapter provides basic information on the following supported switches, host bus adapters, operating system, and arrays.

Switches

This section includes basic information on EMC®-supported switches. For more information, including how to install these switches, refer to the appropriate hardware reference manuals, available on Powerlink. For the most up-to-date information on supported switches, refer to the *EMC Support Matrix* (ESM), available through the E-Lab Interoperability Navigator (ELN), located at http://elabnavigator.EMC.com, under the **PDFs and Guides** tab.

Connectrix B

The EMC Connectrix B series offers the industry's most extensive selection of networked storage connectivity products, ranging from highly available enterprise-class directors to edge switches and management tools. All Connectrix systems concurrently support multiple hosts and operating systems, enabling multiple servers to connect to multiple storage systems from EMC as well as other vendors. This flexible family of platforms helps facilitate the consolidation of information onto fewer storage systems for improved management, protection, and availability.

Figure 1   Connectrix switch

All current EMC Connectrix B series switches running v6.1.0f or higher are supported.

Connectrix MDS

The Connectrix MDS 9000 Family of multilayer directors and fabric switches offers intelligent fabric-switching services that realize maximum performance while ensuring high reliability levels. They combine robust and flexible hardware architecture with multiple layers of network and storage management intelligence. This powerful combination enables highly available, scalable storage

networks that provide advanced security and unified management features. These release notes provide the most up-to-date information on new features and changes.

The Connectrix® MDS 9000 Family provides intelligent networking features such as multiprotocol and multitransport integration, virtual SANs (VSANs), advanced security, sophisticated debug analysis tools, and unified SAN management.

Note: Refer to the *EMC Support Matrix* available through E-Lab™ Interoperability Navigator at http://elabnavigator.EMC.com for the latest information on supported configurations.

The answers to many MDS 9000 related questions can be found in the EMC Knowledgebase located at http://Powerlink.EMC.com

MDS 9000 documentation can be found on www.cisco.com.

QLogic

QLogic Fibre Channel switches, SANbox 5600, 5800, and 9000 offer SAN connectivity to meet the needs of any IT environment. Each QLogic SAN switch comes pre-loaded with all the management tools required to easily connect servers to storage.

For more information on QLogic SANbox switches, refer to the appropriate hardware reference manuals and data sheets, available on the QLogic website, at http://qlogic.com.

Addressing Symmetrix devices

Note: Currently FC-AL is not supported with the MAC OS.

Fabric addressing

Each port on a device attached to a fabric is assigned a unique 64-bit identifier called a World Wide Port Name (WWPN). These names are factory-set on the HBAs in the hosts, and are generated on the Fibre Channel directors in the EMC Symmetrix® system.

When an N_Port (host server or storage device) connects to the fabric, a login process occurs between the N_Port and the F_Port on the fabric switch. During this process, the devices agree on such operating parameters as class of service, flow control rules, and fabric addressing. The N_Port's fabric address is assigned by the switch and sent to the N_Port. This value becomes the Source ID (SID) on the N_Port outbound frames and the Destination ID (DID) on the N_Port inbound frames.

The physical address is a pair of numbers that identify the switch and port, in the format **s**,**p**, where **s** is a domain ID and **p** is a value associated to a physical port in the domain. The physical address of the N_Port can change when a link is moved from one switch port to another switch port. The WWPN of the N_Port, however, does not change. A Name Server in the switch maintains a table of all logged-in devices, so N_Ports can automatically adjust to changes in the fabric address by keying off the WWPN.

The highest level of login that occurs is the process login. This is used to establish connectivity between the upper-level protocols on the nodes. An example is the login process that occurs at the SCSI FCP level between the HBA and the Symmetrix system.

ATTO Host Bus Adapters

For the most up-to-date information on supported HBAs, refer to the *EMC Support Matrix* (ESM), available through the E-Lab Interoperability Navigator (ELN), located at http://elabnavigator.EMC.com, under the **PDFs and Guides** tab.

For more information on various supported HBAs, refer to the EMC page on the ATTO website, at www.attotech.com/solutions/emc.

ATTO host adapters deliver fast performance, one-click installation, and smooth data streaming. Featuring a complete line of SCSI and Fibre Channel host adapters for PCI, PCIe, and PCIe 2.0-based computing environments, ATTO host adapters optimize the link between hardware and software, delivering high performance and interoperability.

The approved ATTO technology HBA options are:

- 4 Gb PCI-e single channel FC HBAs
- 4 Gb PCI-e dual channel FC HBAs
- 4 Gb PCI-e quad channel FC HBAs
- 8 Gb Celerity FC-82EN FC HBA

Notes

Note the following:

- ATTO recommends hard setting speed on the HBA side.
- ATTO has configuration tool to see basic settings of HBA.
- Adding a LUN on the fly is disruptive.

Technical highlights

The following list provides technical highlights:

- Dual-channel 4-Gigabit Fibre Channel host adapter
- 800 MB/s throughput per channel in full-duplex mode
- Exclusive Advanced Data Streaming (ADS) Technology
- Proven industry interoperability with support for SNIA HBA API

- Driver support for Linux, Windows and Mac OS X
- Backward compatible with 2-Gb and 1-Gb legacy Fibre Channel products
- Includes two 4-Gigabit LC SFPs
- Includes high- and low-profile form factor bracket
- RoHS Compliant

ATTO Configuration Tool

The ATTO Configuration Tool helps you customize the settings of your ATTO host adapter to maximize the performance of your storage connection. While the factory settings on your host adapter provide excellent performance for a wide range of applications, some specialized applications may benefit from adjustment of the adapter settings to tune the adapter for a specific performance range. The Configuration Tool provides a powerful tool to help you optimize your host adapter. Additionally, this utility may also be used to verify the drivers and flash files currently in use and update the hardware flash.

For more information on the ATTO Configuration Tool, refer to the EMC page on the ATTO website, at www.attotech.com/solutions/emc.

Manually register Apple server with an HBA

⚠️ **IMPORTANT**

This procedure assumes that all zoning is already in place.

The following steps show how to manually register an Apple server with an HBA.

1. From an EMC-approved switch, issue the **switchshow** command and copy the WWN of your HBA port: 21:00:00:10:86:01:f8:c7 (port 14 in Figure 2).

Figure 2   WWN of HBA port

2. Right-click on your array and then click **Connectivity Status** from Navisphere, as shown in Figure 3.
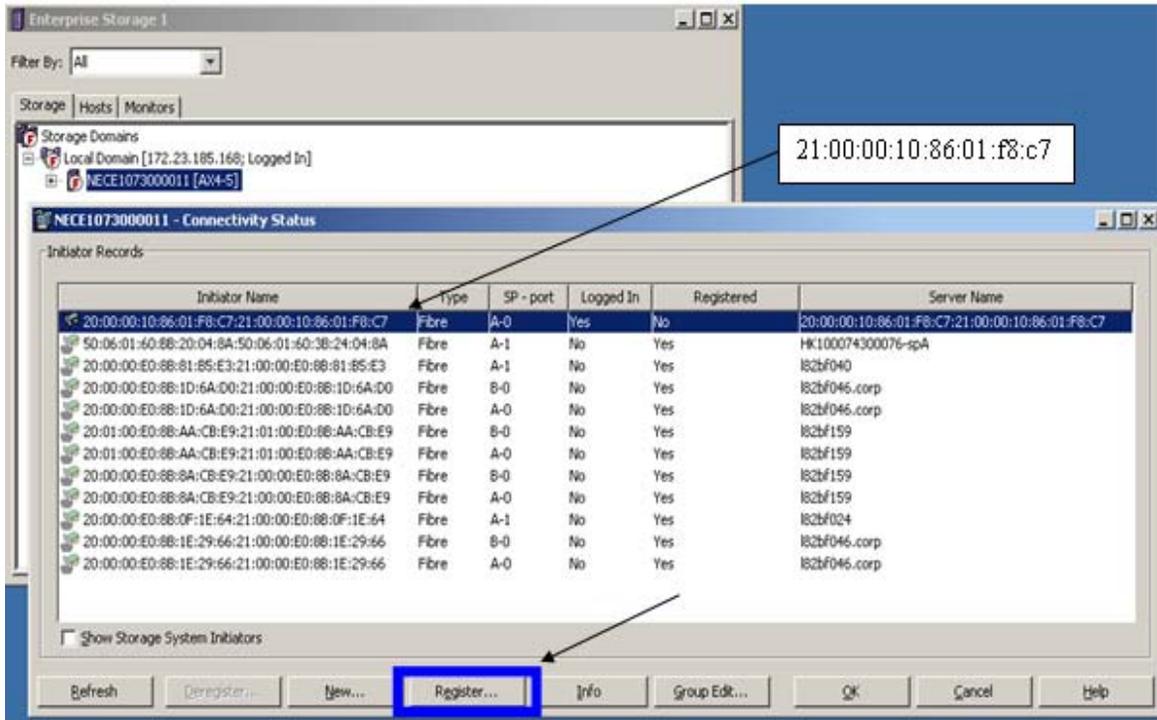
Figure 3   Connectivity Status window

3. Find your HBA.

   As shown in Figure 3, this HBA is logged in, but not registered. Therefore, the HBA needs to be manually registered.

4. Highlight your HBA and click **Register**.

The **Register Initiator Record** dialog box displays, as shown in Figure 4.
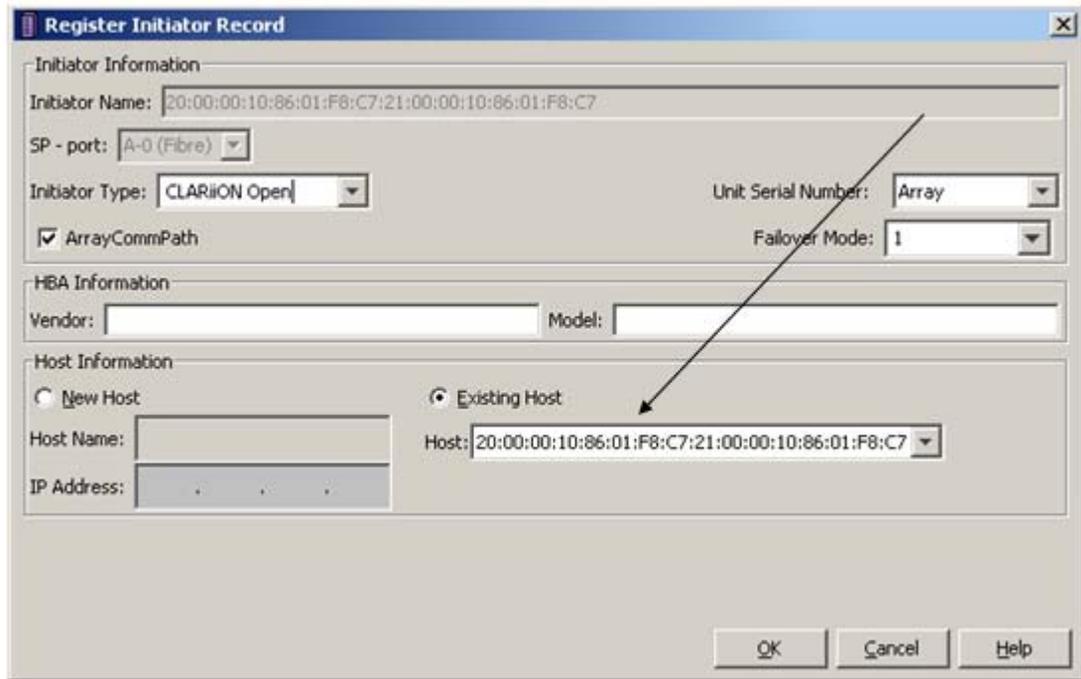
Figure 4   Register Initiator Record dialog box

5. Complete the following in this dialog box:

- In the **Initiator Type** drop-down menu, choose **CLARiiON Open**.
- Check the **ArrayCommPath** box.
- In the **Failover Mode** drop-down menu, choose **4.**
- Click Existing Host.
- Select your HBA from the drop-down menu.
- Click **OK**.

A confirmation box displays, as shown in Figure 5.

Figure 5   Confirm: Register Initiator Record dialog box

6. Click **Yes**.

The **Message: Create Initiator Record** dialog box should appear with a "Success" message, as shown in Figure 6.

Figure 6   Message: Create Initiator Record dialog box

7. Click **OK**.

The **Connectivity Status** window displays. It should now show the Initiator Name as logged in and registered, as shown in Figure 7.

Initiator Name	Type	SP - port	Logged In	Registered	Server Name
20:00:00:10:86:01:F8:C7:21:00:00:10:86:0...	Fibre	A-0	Yes	Yes	20:00:00:10:86:01:F8:C7:21:00:00:10:86:0...
50:06:01:60:BB:20:04:8A:50:06:01:60:3B:2...	Fibre	A-1	No	Yes	HK100074300076-spA
20:00:00:E0:8B:81:B5:E3:21:00:00:E0:8B:8...	Fibre	A-1	No	Yes	I82bf040
20:00:00:E0:8B:1D:6A:D0:21:00:00:E0:8B:1...	Fibre	B-0	No	Yes	I82bf046.corp
20:00:00:E0:8B:1D:6A:D0:21:00:00:E0:8B:1...	Fibre	A-0	No	Yes	I82bf046.corp
20:01:00:E0:8B:AA:CB:E9:21:01:00:E0:8B:A...	Fibre	B-0	No	Yes	I82bf159
20:01:00:E0:8B:AA:CB:E9:21:01:00:E0:8B:A...	Fibre	A-0	No	Yes	I82bf159
20:00:00:E0:8B:8A:CB:E9:21:00:00:E0:8B:8...	Fibre	B-0	No	Yes	I82bf159
20:00:00:E0:8B:8A:CB:E9:21:00:00:E0:8B:8...	Fibre	A-0	No	Yes	I82bf159
20:00:00:E0:8B:0F:1E:64:21:00:00:E0:8B:0...	Fibre	A-1	No	Yes	I82bf024
20:00:00:E0:8B:1E:29:66:21:00:00:E0:8B:1...	Fibre	B-0	No	Yes	I82bf046.corp
20:00:00:E0:8B:1E:29:66:21:00:00:E0:8B:1...	Fibre	A-0	No	Yes	I82bf046.corp

Figure 7   Connectivity Status window

Operating system

Note: Mac OS X Leopard versions 10.5.5 or higher are supported.

One of the many new features in the Mac OS X Leopard is *stacks.* A stack is a Dock item that gives you fast access to a folder of files. When you click a stack, the files within spring from the Dock in a fan or a grid, depending on the number of items (or the preference you set).

Leopard starts you off with two pre-made stacks:

◆ Downloads

The Downloads stack automatically captures files downloaded from Safari, Mail, and iChat.

◆ Documents

The Document stack is helpful for storing presentations, spreadsheets, and word processing files

You can create as many stacks as you wish simply by dragging folders to the right side of your Dock.

Figure 8   Mac OS X Leopard stacks

Arrays

This section briefly discusses the following supported arrays:

- "CLARiiON CX4," next
- "CLARiiON CX3" on page 25
- "Symmetrix DMX-4" on page 26
- "Symmetrix VMAX systems" on page 29

CLARiiON CX4

EMC Fibre Channel CX4 disk-array storage systems, shown in Figure 9, provide terabytes of disk storage capacity, high transfer rates, flexible configurations, and highly available data at low cost. Its hardware RAID features are provided by two storage processors (SPs).

Figure 9   **CLARiiON CX4**

A CX4 system is based on a 4-gigabit disk processor enclosure (called a DPE2). A DPE2 supports a total of 15 Fibre Channel disks within itself and as many as three separate 4-Gb disk-array enclosures (DAE2s) for a total of 60 disks.

Two types of DAE2s are available:

- A standard DAE2 version for high-performance Fibre Channel disks
- A DAE2-ATA version for economical ATA (Advanced Technology Attachment) disks.

CLARiiON CX3

The CX300 2-gigabit disk processor enclosures (DPE2), shown in Figure 10, is an intelligent, highly available, high-performance, high-capacity disk-array storage systems.

Figure 10   CLARiiON CX3

The CX300 uses a Fibre Channel arbitrated loop (FC-AL) or Fibre Channel switch (FC-SW) as its interconnect interface to arrayed disks and to servers that use the storage system.

A CX300 can support as many as three additional disk-array enclosures (DAEs, also called array modules). A DAE is a basic enclosure without a storage processor (SP); it includes either high-performance Fibre Channel or economical ATA disks. The CX300 and three additional DAEs support up to 60 disk modules in a single disk-array storage system. You can place the DAEs in the same cabinet as the CX300, or in one or more separate cabinets.

CX-Series storage systems support both 2-gigabit disk-array enclosures (DAE2s) and 2-gigabit point-to-point disk-array enclosures (DAE2Ps).

The CX300 connects to the external Fibre Channel environment using small form factor (SFF) LC optical transceivers on the storage processor. The CX300i attaches to an external Ethernet environment using standard RJ45 LAN connectors and Ethernet cables.

High-availability features are standard.

Symmetrix DMX-4

The Symmetrix DMX™-4 model, shown in Figure 11 on page 26, establishes a performance and capacity trajectory for the highest of the high-end enterprise systems.

The DMX-4 offers 4 Gb/s front-end and back-end that provides increased performance without increasing power and cooling. The point-to-point Fibre Channel back end has advanced disk isolation capabilities to improve serviceability.

Figure 11   **Symmetrix DMX-4**

The Symmetrix DMX-4 fully leverages the EMC industry-leading storage management functionality and introduces the economic benefits of scalable packaging to the high-end storage market.

The Symmetrix DMX-4 is incrementally scalable, supporting from 96 to 1,920 2 Gb/s or 4 Gb/s high-performance Fibre Channel disk drives and 4 Gb/s SATA II drives, providing a maximum raw capacity of approximately 1 PB.

To support the massive scalability of DMX-4 configurations, the DMX architecture has been expanded and improved to deliver higher throughput (1 GB/s links) and increased I/O performance (four dual 1.3 GHz PPC processor complexes per director).

The field-proven Direct Matrix Architecture provides dedicated, non-blocking interconnects between I/O directors and global memory regions.

Combined with expanded global memory director technology and the dynamically optimized caching algorithms of the Enginuity™ storage operating environment, systems based on the Symmetrix DMX architecture deliver scalable performance to meet the most demanding information.

EMC Symmetrix applications and offerings

This section contains information on Symmetrix applications and offerings. EMC Symmetrix documentation is available on http://Powerlink.EMC.com.

Table 1 lists the EMC Symmetrix arrays supported on the Linux Platform, along with the minimum microcode revisions. The particular code levels supported within a Microcode Family are listed in the Path Management Software table in the *EMC Support Matrix*.

Please refer to the Base Connectivity and the iSCSI Connectivity tables in the *EMC Support Matrix* for supported HBAs, operating system revisions, and servers.

Table 1   Supported EMC Symmetrix arrays

Storage Arrays	Array Code Requirements
Symmetrix DMX-4 [Fibre Channel and iSCSI]	5772 /5773 Microcode Family
Symmetrix DMX-3 [Fibre Channel and iSCSI]	5771/5772 Microcode Family
Symmetrix DMX/DMX-2 [Fibre Channel and iSCSI]	5670/5671 Microcode Family
Symmetrix 8000 Series	5567 and higher Microcode Family

Symmetrix array-specific settings   When attaching an Apple host to a Symmetrix system, use the default FA settings listed in Table 2.

Table 2   Linux default FA settings (1 of 2)

Bits to Set	Description
C	Common serial number
SPC-2	SCSI Primary Commands - 2
EAN	Enable Auto Negotiation
PP (FC-SW only)	Enable point-to-point
UWN	Unique WorldWide Name

Table 2   Linux default FA settings (2 of 2)

Bits to Set	Description
D (2.6 kernel)	Disable queue reset on Unit Attention
SC3	SCSI-3 compliant
ACLX	Access Logix (for VMAX only)

Symmetrix SPC-2 director bit considerations

Enginuity™ code versions 5671.58.64 (and later) for DMX and DMX-2, Enginuity code versions 5771.87.95 (and later) for DMX-3, and Enginuity code versions for DMX-4 and VMAX™, provide support for compliance with newer SCSI protocol specifications; specifically, SCSI Primary Commands - 2 (SPC-2) as defined in the SCSI document at http://www.t10.org/ftp/t10/drafts/spc2/spc2r20.pdf.

The SPC-2 implementation in Enginuity includes functionality which, based on OS and application support, may enhance disk-attach behavior to use newer SCSI commands optimized for a SAN environment (as implemented in Fibre Channel), as opposed to legacy (non SPC-2) functionality, which was targeted for older SCSI implementations utilizing physical SCSI bus-based connectivity (which cannot leverage the enhanced functionality of newer SCSI specifications).

In environments sharing director ports between hosts with multiple vendor operating systems, ensure that all hosts' operating systems are capable of supporting the SPC-2 functionality before enabling it on the port. If any OS sharing the affected director port does not support SPC-2 functionality, the SPC-2 bit cannot be set on a per-port basis and must be set on a per-initiator basis using Solutions Enabler 6.4 or later CLI. Refer to the *EMC Solutions Enabler Symmetrix Array Management CLI Product Guide*, available on Powerlink, for details regarding how to set the SPC-2 bit on a per-initiator basis.

SPC-2 must be enabled for all initiators on a per-host basis, globally, so if SPC-2 conformance is enabled for a specific Symmetrix device visible to a specific host, SPC-2 conformance must be enabled for all paths to that same device and from that same host.

Unless otherwise specified in the *EMC Support Matrix*, SPC-2 conformance is supported for new Symmetrix devices only. Migrations from SPC-2 disabled to SPC-2 enabled configurations, either online or offline, are not currently supported.

Symmetrix VMAX systems

EMC Symmetrix VMAX systems are enterprise-class storage platforms intended for open systems and mainframe computing.

Symmetrix VMAX systems run Enginuity 5874 and greater, and are offered in two models:

◆ The highly scalable Symmetrix VMAX system

◆ The single engine Symmetrix VMAX SE system

The VMAX Engine is a system bay component that provides physical memory, front-end host connectivity (including SRDF), back-end connectivity, and connection to other VMAX Engines. Each VMAX Engine is fully populated, provides 32, 64, or 128 GB of physical memory, multiple host configuration options, and connection to eight disk array enclosures.

Figure 12 shows the VMAX Engine.
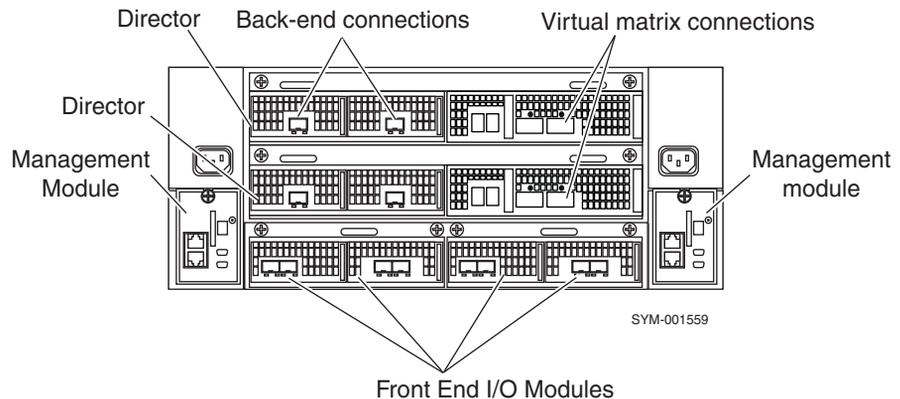
Figure 12   VMAX Engine

Symmetrix VMAX systems provide a maximum of eight VMAX Engines within a system bay. Symmetrix VMAX SE systems provide one VMAX Engine within a system bay. Each VMAX Engine includes:

◆ Two directors that support front-end, back-end and SRDF connections. Each director has:

• 32, 64, or 128 GB of physical memory.

- • One serial I/O interface board (SIB) that connects the director and the Matrix Interface Board Enclosure (MIBE).
- • Two Back End I/O Modules (4 ports, 4-Gb/s) that connect to storage bay drives.

◆ Two I/ O Module carriers that provide connectivity between the director and the front-end and (optional) SRDF I/O ports. Front End I/O Modules support:

- • Fibre Channel host connectivity (4 ports, 2 or 4 Gb/s).
- • Fibre Channel SRDF connectivity.

> **Note:** For the most up-to-date support information, refer to the *EMC Symmetrix VMAX Series Product Guide*.

- • FICON host connectivity (2 ports, 2 or 4 Gb/s).
- • iSCSI host connectivity (2 ports, 1 Gb/s).
- • GigE SRDF connectivity.

> **Note:** For the most up-to-date support information, refer to the *EMC Symmetrix VMAX Series Product Guide*.

◆ Two management modules that provide environmental monitoring.

◆ Two enclosure standby power supplies.

◆ Four cooling fans.

For more information on VMAX and VMAX Enginuity features, consult the *EMC Symmetrix VMAX Series Product Guide*. For details on SRDF-related I/O module configurations, please refer to the "I/O Module Configurations" and the "Typical Configurations" sections of the *EMC Symmetrix VMAX Series Product Guide*.

SCSI-3 FCP addressing

The Symmetrix Fibre Channel director extracts the SCSI Command Descriptor Blocks (CDB) from the frames received through the Fibre Channel link. Standard SCSI-3 protocol is used to determine the addressing mode and to address specific devices.

Apple and the Symmetrix system support Volume Set addressing as defined by the SCSI-3 Controller Commands (SCC). Volume Set Addressing uses the first two bytes (0 and 1) of the 8-byte LUN

Supported Environments

addressing structure. (Bits 7 and 6 of byte 0 are **01**, identifying the addressing mode as Volume Set.) The remaining six bytes are set to zeros.

The Symmetrix port identifies itself as an array controller in response to a host Inquiry command sent to LUN 00. This identification is done by returning the byte 0x0C in the **Peripheral Device Type** field of the returned data for Inquiry. If the Symmetrix system returns the byte 0x00 in the first byte of the returned data for Inquiry, the Symmetrix system is identified as a *direct access device.*

Upon identifying the Symmetrix system as an array controller device, the host should issue a SCSI-3 Report LUNS command (0xA0), in order to discover the LUNs.

Zoning recommendation

When using Apple hosts in a fabric environment, the recommended zoning methodology is single-initiator zoning. A single-initiator zone consists of only one host bus adapter port. While multiple array target ports may be part of the same zone, it is required that a single zone should not contain target ports from multiple arrays.

Note: Multi-initiator zones are not supported in a Apple fabric environment.

2

Virtual Provisioning

This chapter provides information about Virtual Provisioning and Apple.

Note: For further information regarding the correct implementation of Virtual Provisioning, refer to the *Symmetrix Virtual Provisioning Implementation and Best Practices Technical Note*, available on http://Powerlink.EMC.com.

Virtual Provisioning on Symmetrix

Virtual Provisioning enables organizations to improve speed and ease of use, enhance performance, and increase capacity utilization for certain applications and workloads. EMC Symmetrix Virtual Provisioning integrates with existing device management, replication, and management tools, enabling customers to easily build Virtual Provisioning into their existing storage management processes.

Virtual Provisioning, which marks a significant advancement over technologies commonly known in the industry as "thin provisioning," adds a new dimension to tiered storage in the array, without disrupting organizational processes.
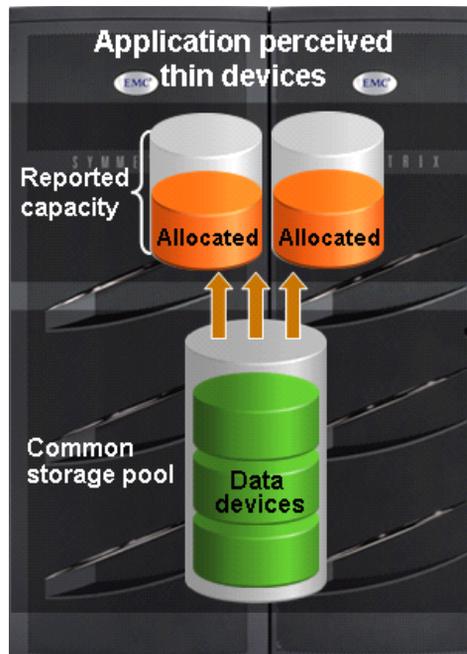
Figure 13   Virtual Provisioning on Symmetrix

Terminology

This section provides common terminology and definitions for Symmetrix and thin provisioning.

Symmetrix   Basic Symmetrix terms include:

Device	A logical unit of storage defined within an array.
Device Capacity	The storage capacity of a 'Device'.
Device Extent	Specifies a quantum of logically contiguous blocks of storage.
Host Accessible Device	A device that can be made available for host use.
Internal Device	A device used for a Symmetrix internal function that cannot be made accessible to a host.
Storage Pool	A collection of 'Internal Devices' for some specific purpose.

Thin provisioning   Basic thin provisioning terms include:

Data Device	An 'Internal Device' that provides storage capacity to be used by 'Thin Devices'.
Thin Pool	A collection of 'Data Devices' that provide storage capacity for 'Thin Devices'.
Thin Pool Capacity	The sum of the capacities of the member 'Data Devices'.
Thin Pool Allocated Capacity	A subset of 'Thin Pool Enabled Capacity' that has been allocated for the exclusive use of all 'Thin Devices' bound to that 'Thin Pool'.
Bind	Refers to the act of associating one or more 'Thin Devices' with a 'Thin Pool'.

Pre-Provisioning   An approach sometimes used to reduce the operational impact of provisioning storage. The approach consists of satisfying provisioning operations with larger devices that needed initially, so that the future cycles of the storage provisioning process can be deferred or avoided.

Over-Subscribed Thin Pool   A thin pool whose thin pool capacity is less than the sum of the reported sizes of the thin devices using the pool.

Data Device Extent   The minimum quantum of storage that is allocated at a time when dedicating storage from a thin pool for use with a specific thin device.

3

iSCSI Initiator

This chapter provides the following information for ATTO's Xtend SAN iSCSI Initiator for Mac OS X:

:

iSCSI Initiator overview

Internet SCSI (iSCSI) enables key applications including collaborative digital video and audio workflows, laptop connectivity to SANs, and remote backups. Users can take advantage of their existing Ethernet infrastructure while gaining the benefits of network attached storage, avoiding both the cost and complexity associated with Fibre Channel.

Note: For the most up-to-date information, always consult the *EMC Support Matrix* (ESM), available through E-Lab Interoperability Navigator (ELN) at http://elabnavigator.EMC.com, under the **PDFs and Guides** tab.

Benefits

Benefits include:

- High performance
  - Transports more data in less time
  - Supports multiple streams of uncompressed video
- iSCSI error handling and recovery
- Intuitive GI-based installation and management
- Compatible with
  - Mac OS xv10.4.x - 10.5.x
  - Leading Ethernet switches, network interface cards, and iSCSI storage
  - Leading ISV solutions
  - All ATTO iPBridges and leading iSCSI targets
- Supports
  - CHAP
  - iSNS Client
  - Login redirect functionality
  - Latest iSCSI specifications
  - Header and data digest
- Enhanced data protection
  - Copies data to multiple remote sites
- One-click installation

Installing the ATTO Xtend SAN iSCSI Initiator

Note: Refer to www.attotech.com for more information on ATTO Host Bus Adapters or the ATTO Configuration Tool.

A web page dedicated to EMC customers, including information on various supported drivers and download information, is available at www.attotech.com/solutions/emc.

Management and configuration are simplified through an intuitive GUI, which guides users through the initial setup.Its simplicity allows even non-technical users to set up and manage a server with just a few mouse clicks.

To install the Xtend SAN iSCSI Initiator, follow these steps:

1. Install the Xtend SAN CD, shown in Figure 14.

Figure 14   ATTO Xtend SAN MAC iSCSI Initiator CD

An **iSCSI Xtend SAN** icon, shown in Figure 15, is placed at the location you chose to install the initiator.

Figure 15   Xtend SAN icon

2. Double-click the **iSCSI Xtend SAN** icon.

The **Initiator** dialog box displays, as shown in Figure 16 on page 40. A SAN directory is created, as shown in the left panel.
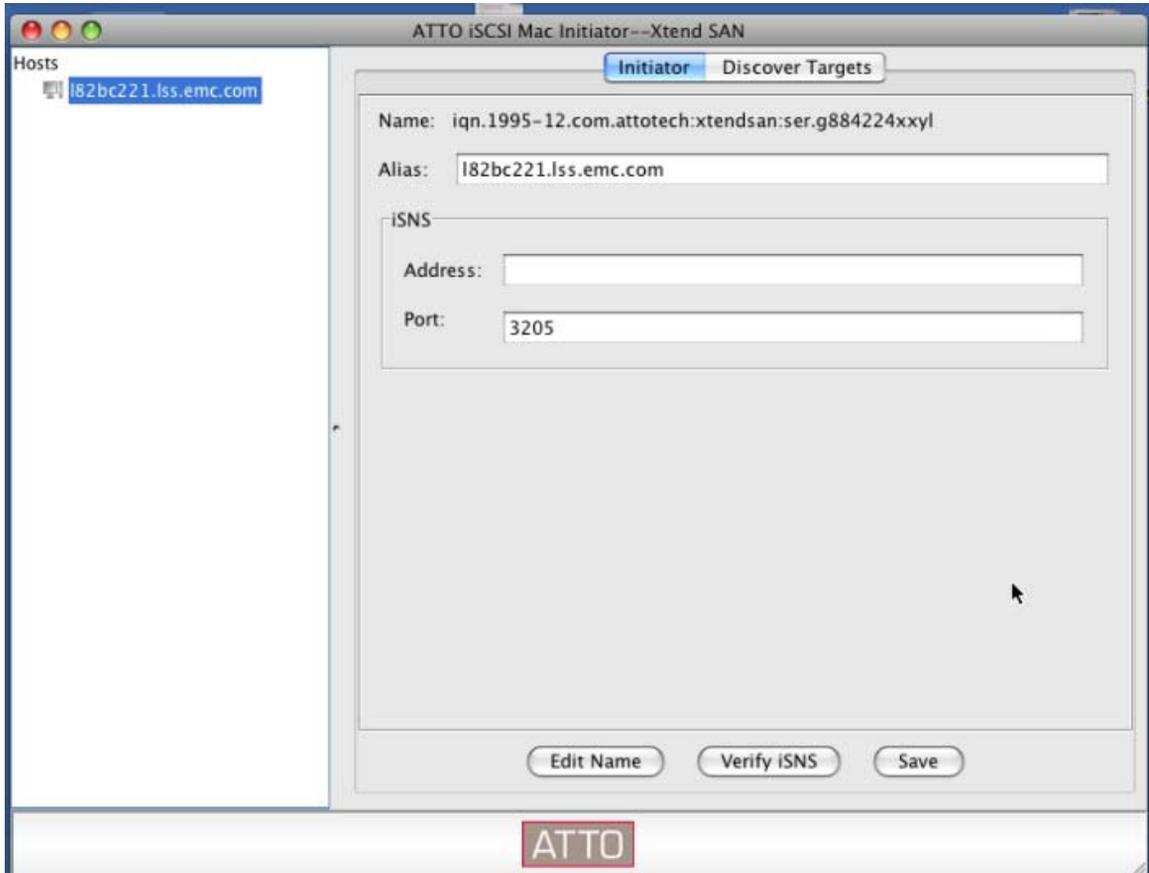
Figure 16   Initiator dialog box showing Xtend SAN host directory

3. Click **Discover Targets**.

The **Discover Targets** dialog box appears, as shown in Figure 17.
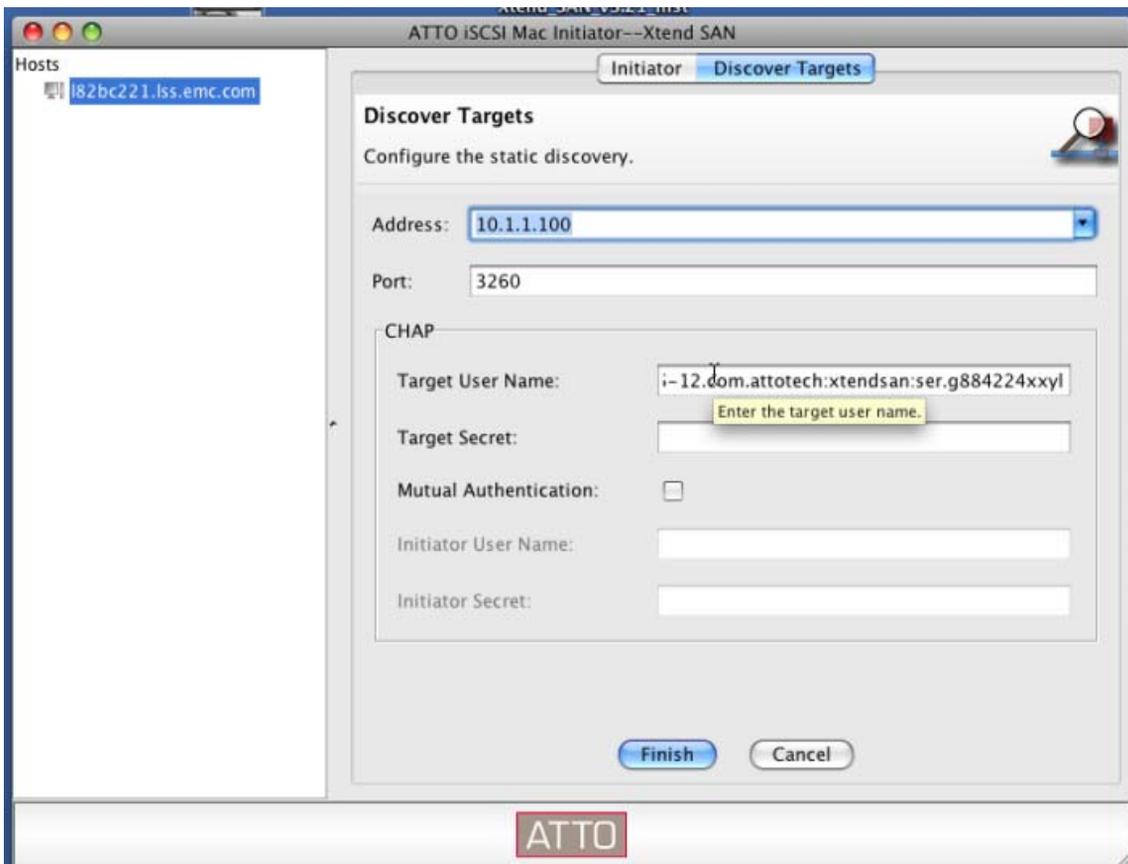
Figure 17   Discover Targets dialog box

4. Type the target address.

In this example, **10.1.1.100** is used.

The **Setup** window appears, displaying the iqn of the Celerra target discovered, as shown in the left panel in Figure 18.
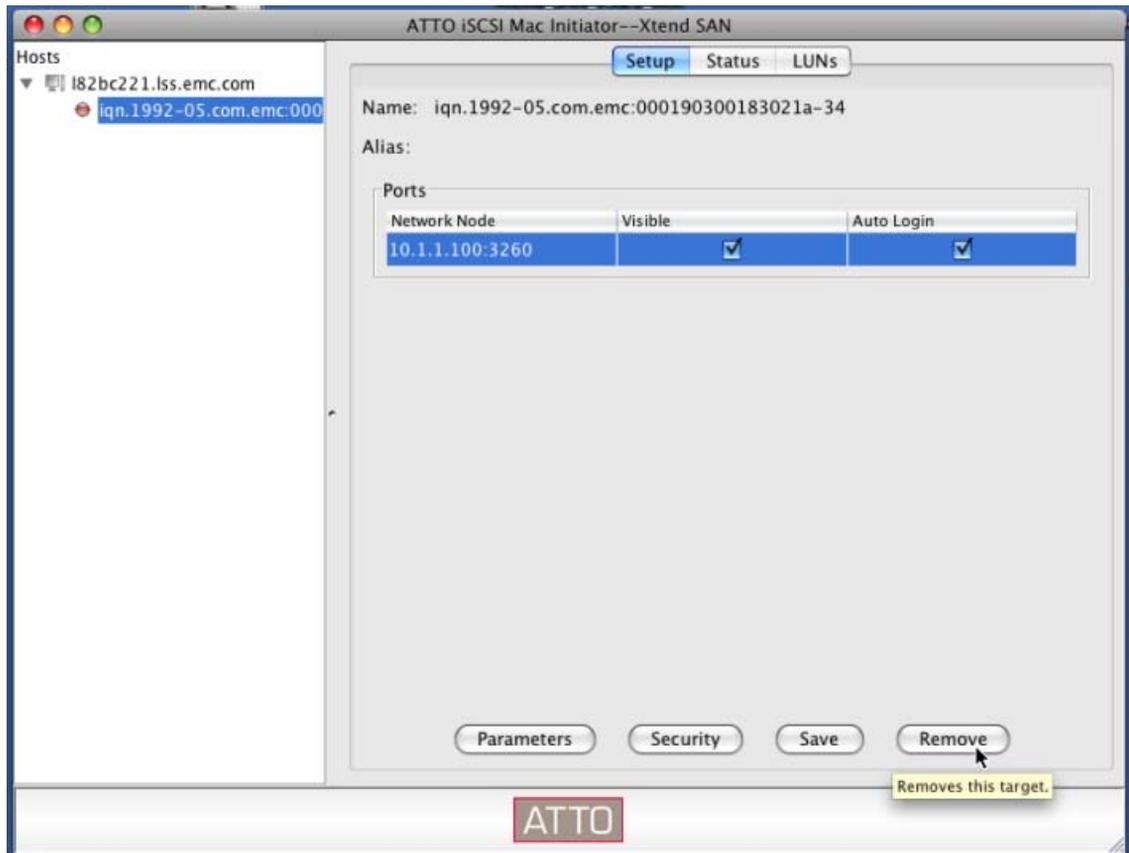
Figure 18   Setup dialog box

5. Enable **Visible** and **Auto Login** by clicking the checkbox, then click anywhere on the line to highlight the IP address, as shown in Figure 18.

6. Click **Parameters**.

The **Configure Parameters** dialog box displays, shown in
Figure 19
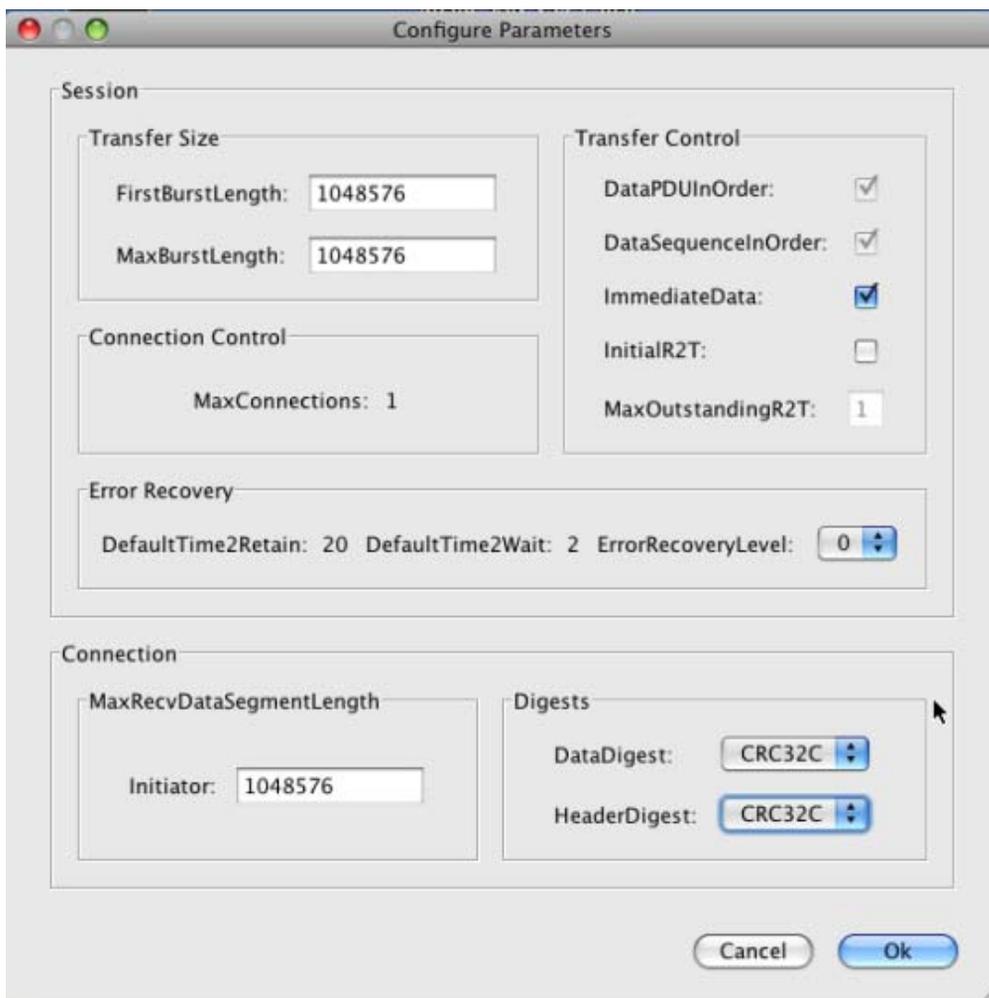
Figure 19   Configure Parameters dialog box

7. (Optional): Choose the DataDigest and/or HeaderDigest.

8. Click **OK**.

    The **Setup** dialog box (shown in Figure 18 on page 42) displays
    again.

9. Click **Save** in the **Setup** dialog box, as shown in Figure 18.

10. Click the **Status** tab.

The **Status** dialog box displays, as shown in Figure 20.

11. Click **Login.**

An **Operation in progress** bar display, as shown in Figure 20.
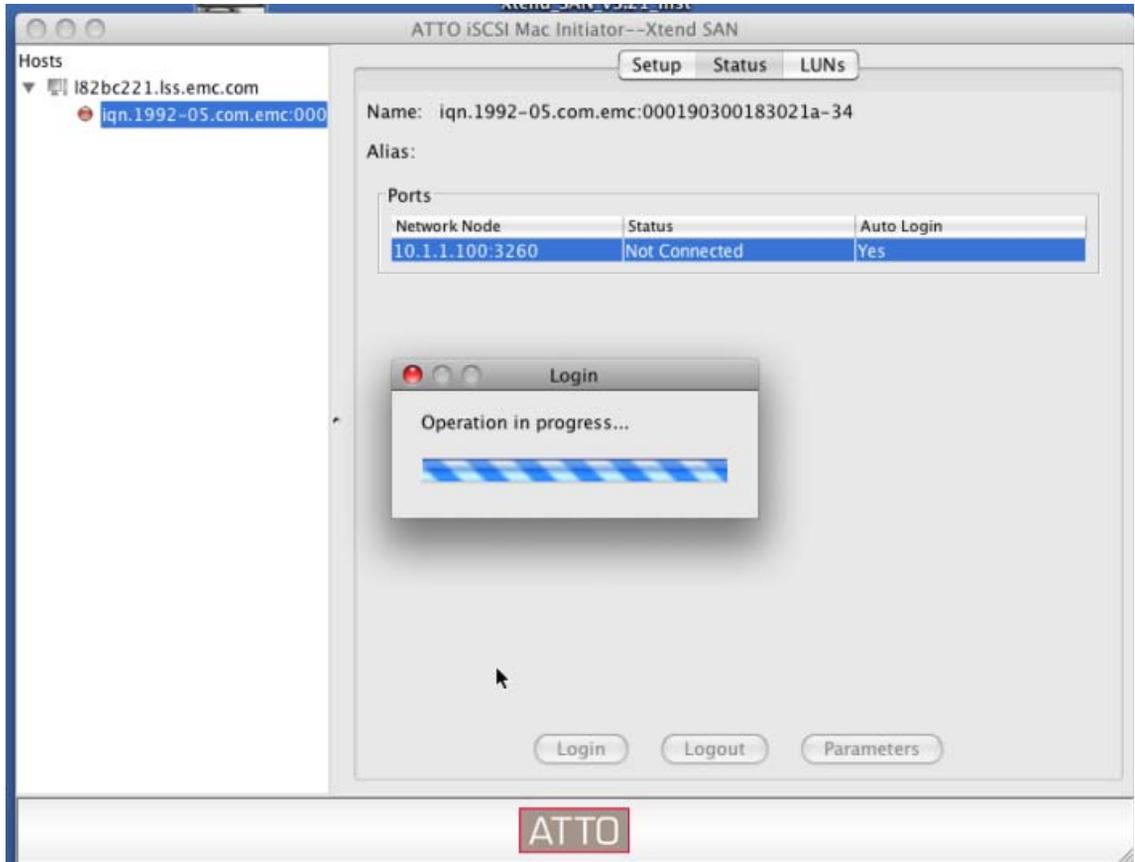
Figure 20   Status dialog box

12. When the operation is complete, the red icon beside the iqn name in the left panel turns green, as shown in Figure 21.
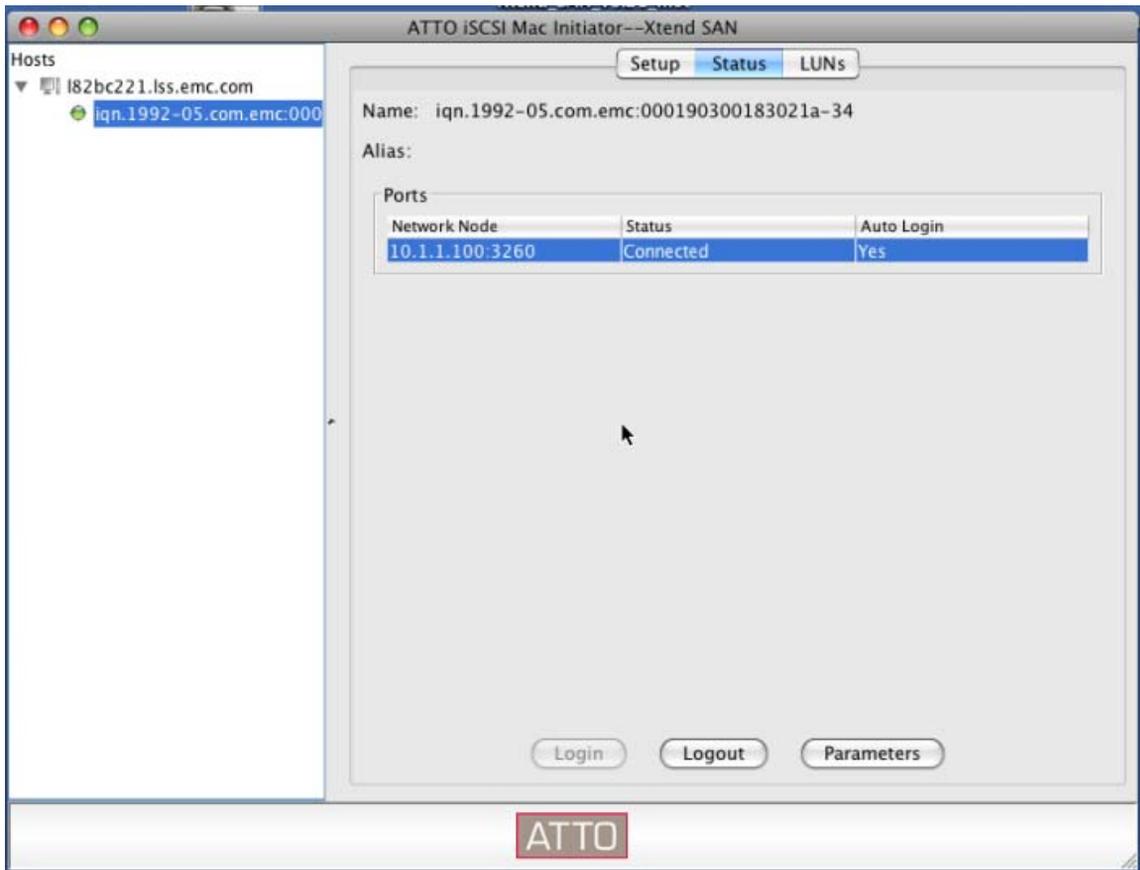
Figure 21   Green icon beside iqn in left panel shows connection

13. To verify the connection, click **LUNs**.

The **LUNs** dialog box displays, as shown in Figure 22. If the initiator is connected to the Celerra, the LUNs display.
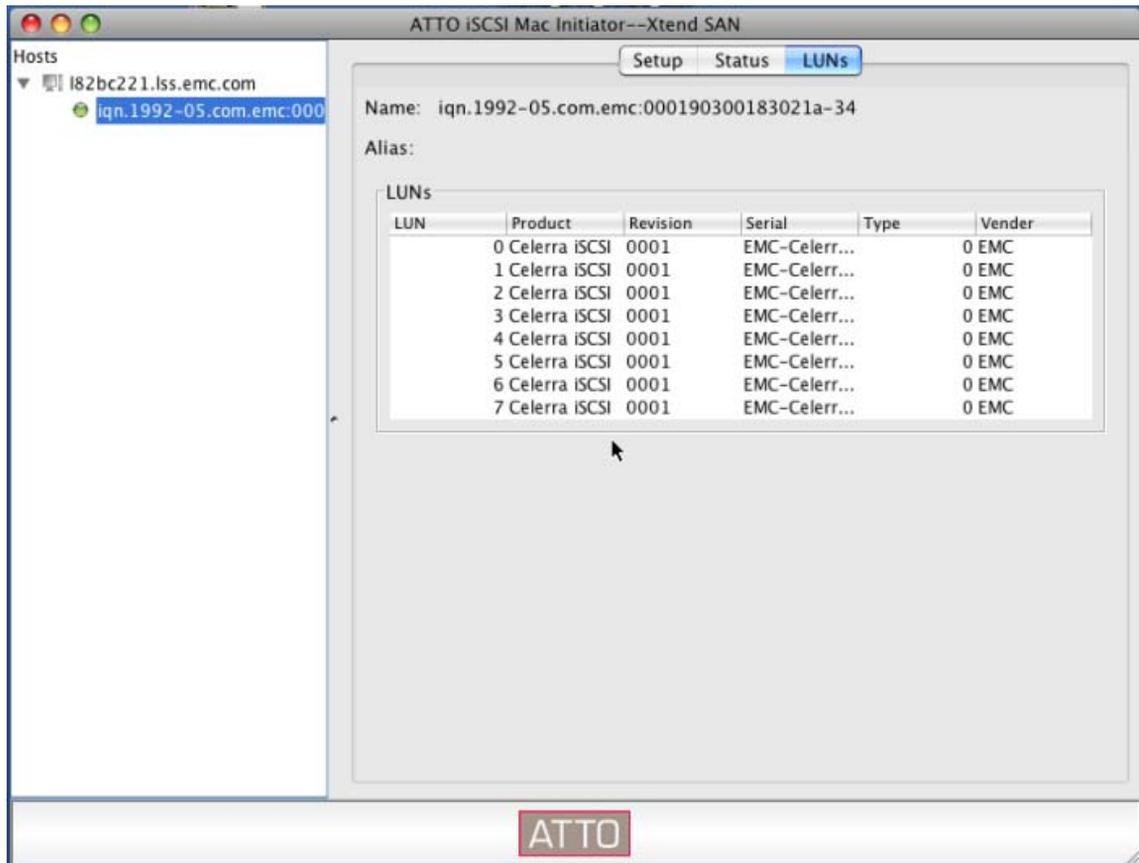
Figure 22   LUNs dialog box

14. To see all the Celerra LUNs available, go to **Disk Utilities**.

15. The **EMC Celerra iSCSI Media** dialog box displays, as shown in Figure 23.
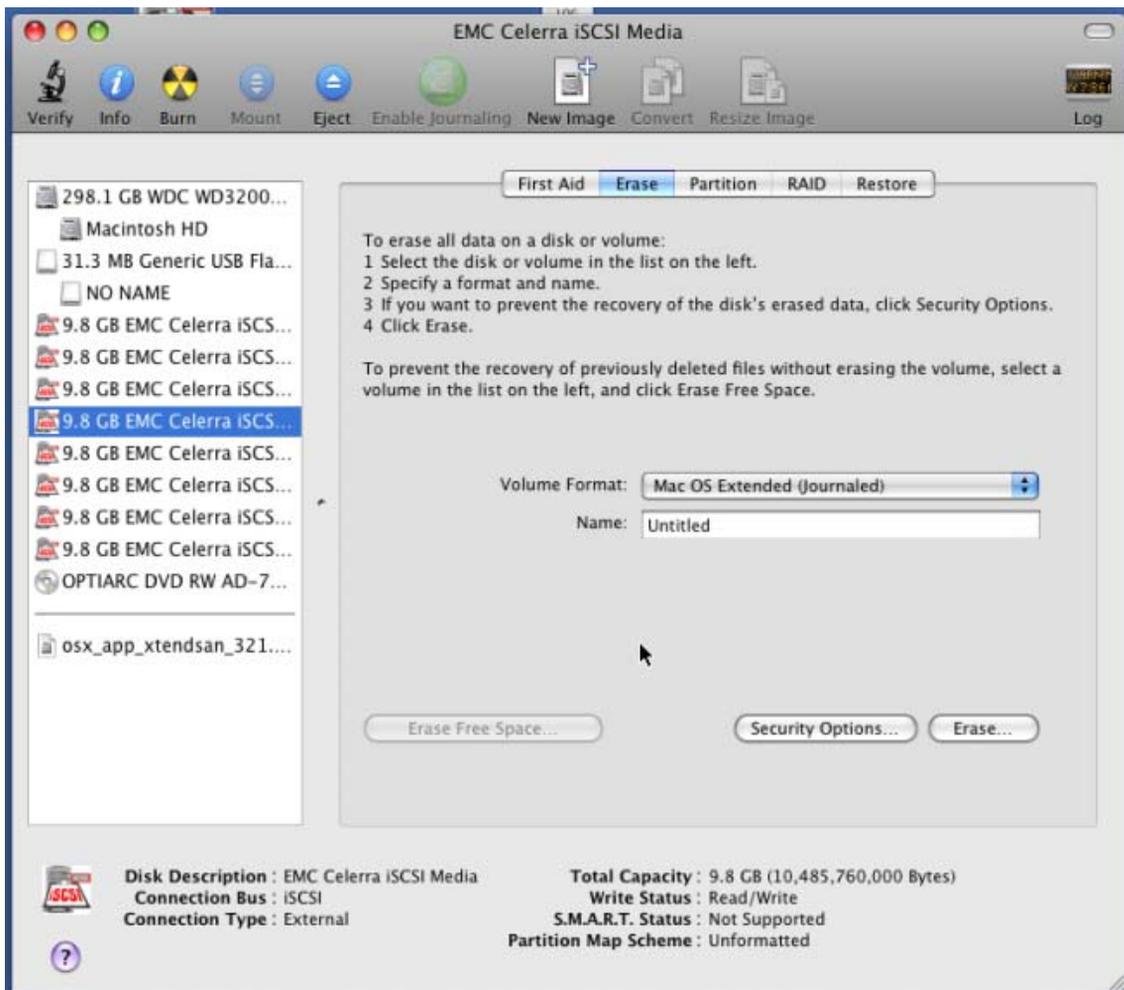
Figure 23   EMC Celerra iSCSI Media dialog box

16. Choose the LUNs you want to use.

    The installation is complete and you can begin using the application.

17. To leave the application, click the red icon on the top left of the dialog box.

Supported configurations

E-Lab has qualified the ATTO Xtend SAN iSCSI Initiator with EMC Celerra NAS and CLARiiON systems.

Figure 24 shows the E-Lab testing environment for the ATTO Xtend SAN iSCSI Initiator with the Celerra NAS system and Symmetrix/CLARiiON storage systems.
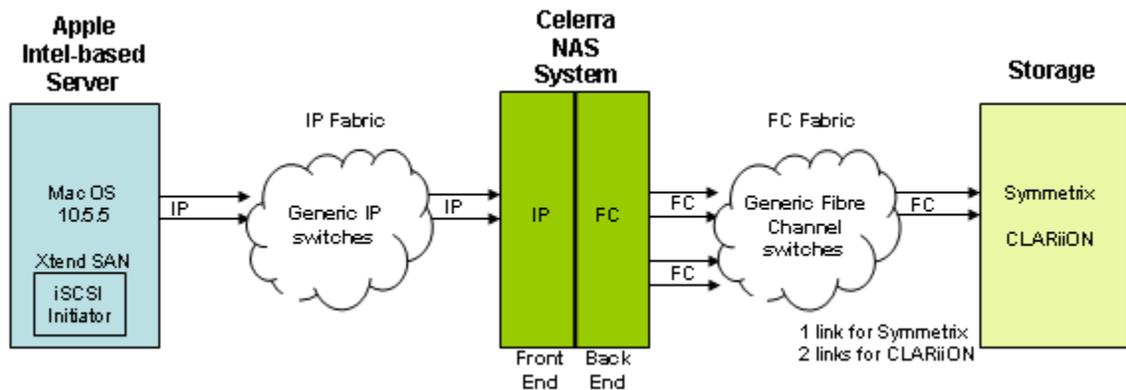

Figure 24   E-Lab testing environment for Celerra NAS system

Figure 25 shows the E-Lab testing environment for the ATTO Xtend SAN iSCSI Initiator with Symmetrix and CLARiiON storage systems.
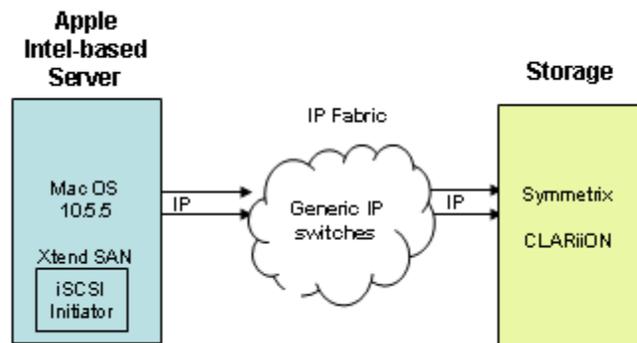
Figure 25   E-lab testing environment for CLARiiON storage system

4

Connectivity in a SAN Environment

This chapter provides an overview of various operating systems, HBAs, and storage-imposed limits when connecting to a SAN environment. In addition, best practices on configuring the Apple host, along with information on system behavior, available features, and useful utilities are discussed.

Operating system limits and guidelines

This section provides operating system limits and restrictions imposed in a SAN environment. Factors such as number of supported host bus adapters, scalability of targets, file system and volume management limits are detailed.

Host initiators

The host initiators may be single or dual channel HBAs. The number of host initiator ports on a server is also limited by the number of HBA slots available on the server and supported by the server vendor.

Note: EMC does not support the mixing of HBAs from different vendors or mixing HBAs with different PCI-interfaces on the same host.

Devices and operations

This section provides an overview of mechanisms provided by an Apple operating system for addressing and utilizing SCSI devices.

SCSI device operation interfaces

Apple provides various device operation interfaces. This includes block and character devices, and raw device interfaces.

The /dev directory contains the files that represent devices attached to the system:

- ◆ /dev/rdisk[0-n] list the raw devices connected to the MAC OS X
- ◆ /dev/rdisk[0-n][0-n] represent the partition number on the raw disk
- ◆ /dev/disk[0-n] contain disk device information
- ◆ /dev/disk[0-n]s[0-n] list the disk partitions

Device details can be found by using the **system_profiler** command found under the Application Folder in MAC OS X.

Following is an example of FC target devices displayed using the system profiler:

```
SCSI Target Device @ 0:

        Node World Wide Name: 50:06:01:60:80:60:04:1B
        Port World Wide Name: 50:06:01:69:00:60:04:1B
        Address Identifier: 1F:61:00
        SCSI Target Identifier: 0
        SCSI Peripheral Device Type: 0
        Manufacturer: DGC
        Model: LUNZ
        Revision: 0219

          SCSI Logical Unit @ 0:

            SCSI Logical Unit Number: 0
            Manufacturer: DGC
            Model: LUNZ
            Revision: 0219

      SCSI Target Device @ 1:

        Node World Wide Name: 50:06:04:8A:CB:37:C3:F3
        Port World Wide Name: 50:06:04:8A:CB:37:C3:F3
```

```
Address Identifier: 36:61:00
SCSI Target Identifier: 1
SCSI Peripheral Device Type: 0
Manufacturer: EMC
Model: SYMMETRIX
Revision: 5671

  SCSI Logical Unit @ 0:

    Capacity: 8.43 GB
    SCSI Logical Unit Number: 0
    Manufacturer: EMC
    Model: SYMMETRIX
    Revision: 5671
    Removable Media: Yes
    Detachable Drive: No
    BSD Name: disk2
    Mac OS 9 Drivers: No
    Partition Map Type: Unknown
    S.M.A.R.T. status: Not Supported
```

LUNZ has been implemented on CLARiiON arrays to make arrays visible to the host OS and PowerPath when no LUNs are bound on that array. When using a direct-connect configuration, and there is no Navisphere Management station to talk directly to the array over IP, the LUNZ can be used as a pathway for Navisphere CLI to send Bind commands to the array.

LUNZ also makes arrays visible to the host OS and PowerPath when the host's initiators have not yet logged in to the Storage Group created for the host. Without LUNZ, there would be no device on the host for Navisphere Agent to push the initiator record through to the array.

⚠ **IMPORTANT**

It is mandatory for the host to log in to the Storage Group.

Once this initiator push is complete, the host will display as an available host to add to the Storage Group in Navisphere Manager (Navisphere Express).

LUNZ should disappear once a LUN zero is bound, or when the Storage Group access has been attained.

Note: For more information refer EMC Knowledgebase Article emc65060.

The disk utility (diskutil) found in the Applications folder is also useful in determining additional information about the devices attached to the MAC OS X.

Multipath

Consult the appropriate HBA vendor for multipath feature support information.

File system and logical volume management

At the time this document was released, EMC only supports the Apple Xsan file system. Refer to Chapter 5, "Xsan," for more information.

5

Xsan

This chapter provides the following information on Xsan:

Overview

Apple's Xsan is a 64-bit cluster file system specifically designed for small and large computing environments that need a high level of data availability. This specialized technology enables multiple Mac desktop and Xserve systems to share RAID storage volumes over a high-speed Fibre Channel network. Each client can read and write directly to the centralized file system, accelerating user productivity while improving workgroup collaboration.

Primary focus of testing

The following lists the primary focus of EMC E-Lab's testing:

- Basic Xsan mode of operation
- FS data availability
- Robust recovery in error conditions
- FS expansion, configurations
- MDC failover with no I/O disruption
- Ease of management

Supported SAN configuration

Figure 26 shows the supported SAN configuration for CLARiiON CX storage systems with FC data ports.
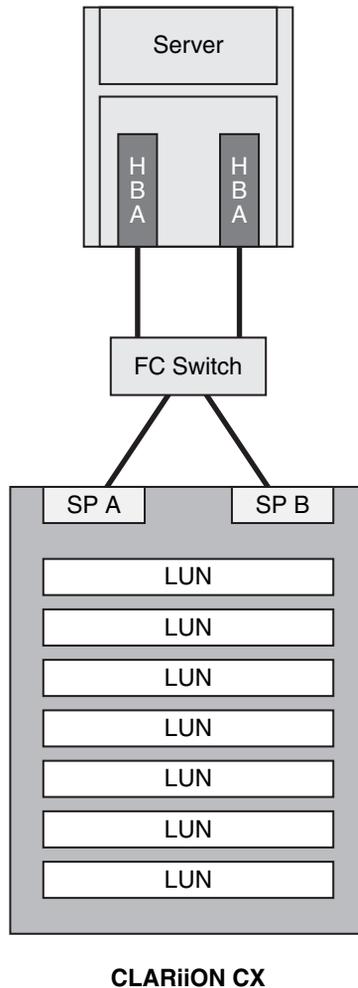
CLARiiON CX

SYM-002061

Figure 26   FC switch connections between server and CLARiiON CX

This configuration has 2 HBAs and 2 SPs.

Supported topology

The Xsan 2.1.1 testing topology is shown in Figure 27.
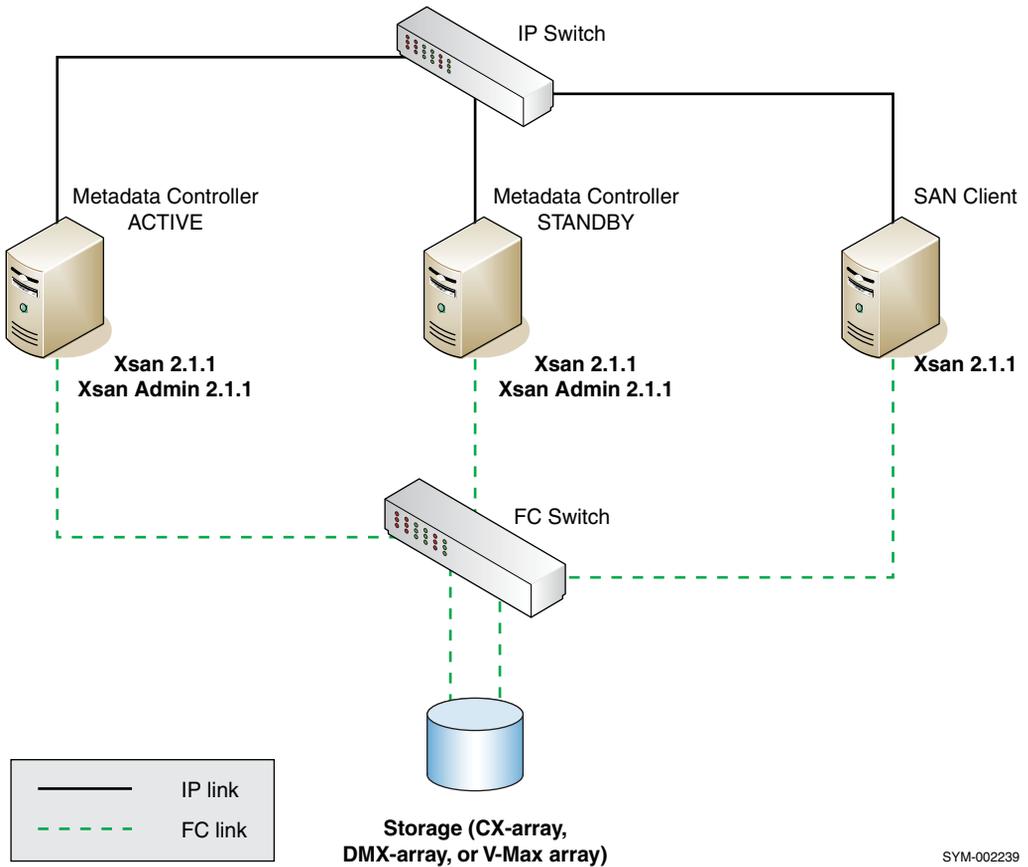
Figure 27   Xsan 2.1.1 CX-array testing topology

Support restrictions

This section lists what EMC supports with limitations and what EMC does not currently support.

Limited support

At the time of this publication, the following have limited support:

◆ No IP client support

◆ Tuning and quota management as per Xsan guide

◆ Support applies only for Xsan admin version 2.1.1

Not supported

At the time of this publication, the following are not supported:

◆ Boot from SAN

◆ CLARiiON layered applications

Index

EMC Host Connectivity Guide for Apple/Mac OS X