



Legacy Technologies

Reference Manual

P/N 300-011-727
REV A02

EMC Corporation

Corporate Headquarters:

Hopkinton, MA 01748-9103

1-508-435-1000

www.EMC.com

Copyright © 2001 – 2011 EMC Corporation. All rights reserved.

Published April, 2011

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date regulatory document for your product line, go to the Technical Documentation and Advisories section on EMC Powerlink.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Preface	9
Chapter 1	Fibre Channel Arbitrated Loop (FC-AL)
FC-AL	16
Overview.....	16
Loop construction.....	18
Loops with hubs	19
Arbitration	20
Symmetrix and Fibre Channel connectivity.....	22
Overview.....	22
Dual port devices.....	23
Arbitrated loop addressing	25
Overview.....	25
Loop ID	26
Primitive signals and sequences	28
Loop Port State Machine (LPSM)	30
Loop initialization.....	33
Overview.....	33
Loop initialization steps	36
Login process.....	40
Arbitration process	43
Overview.....	43
Arbitration process steps.....	44
Access fairness	48
Opening and Closing the loop.....	49
Alternate Buffer-to-Buffer Credit (BB_Credit).....	53

Chapter 2	QuickLoop	
	Overview.....	56
	Need for QuickLoop.....	57
	Defining QuickLoop.....	58
	Concepts, terms, and rules.....	58
	Configuration examples.....	60
	Single QuickLoop configuration.....	60
	Dual QuickLoop configuration.....	61
	Mixed-mode configuration: Fabric and QuickLoop.....	61
	References.....	62
Chapter 3	Bridges	
	Overview.....	64
	SCSI-to-Fibre Channel bridges.....	65
	Operating modes.....	65
	Supported SCSI-to-Fibre Channel fabric bridges.....	66
	SCSI-to-fabric configuration envelope.....	68
	Crosspoint 4200 SCSI-to-fabric configuration.....	69
	Initial setup.....	69
	Target connections.....	70
	Bridge-to-SAN connections.....	70
	Host connections and configuration.....	71
	Fabric zoning.....	71
	Final setup.....	71
	ADIC SAN Gateway SCSI-to-fabric configuration.....	72
	Initial setup.....	72
	Target connections.....	73
	Bridge-to-SAN connections.....	74
	Host connections and configuration.....	74
	Fabric zoning.....	74
	Final setup.....	74
	ADIC SAN Gateway loop-to-fabric configuration.....	75
	Initial setup.....	75
	Target connections.....	76
	Bridge-to-SAN connections.....	77
	Host connections and configuration.....	77
	Fabric zoning.....	77
	Final setup.....	77
	Reference.....	78

Chapter 4	Interfacing Arbitrated Loop to Switched Fabric	
	Overview	80
	Operating modes.....	81
	Storage mode	81
	Host mode.....	81
	Connectivity devices that support FC-AL.....	82
	Connectrix DS-16B, DS-16B2 (Brocade SilkWorM Series) ...	82
	Brocade M Series ES-1000	85
	ADIC SAN Gateway.....	87
	Interfacing arbitrated loop to switched fabric summary	89
	Loop-to-fabric configuration envelope	90
Chapter 5	Storage Area Network Management	
	Distance topology.....	92
	Capacity topology in the loop environment	93
	Consolidation topology in the arbitrated loop environment.....	94
	Combined topologies.....	95
Chapter 6	CNT (Inrange)	
	Configuring CNT (Inrange).....	98
	Supported product.....	98
	Topology support	98
	IOCP considerations	99
Chapter 7	Security Appliances	
	Overview	102
	Decru DataFort FC-Series security appliance	104
	Decru virtualization.....	105
	Decru mapping for the encrypted storage	107
	Decru Cryptainers vault.....	108
	Neoscale CryptoStore security appliance.....	118
	Neoscale CryptoStor FC-2002 for Disk	118
	Neoscale CryptoStor FC702/704 for Tape	121
Glossary		125

	Title	Page
1	Arbitrated loop example	16
2	Arbitrated loop	17
3	Arbitrated loop with 4 nodes	18
4	Arbitrated loop with a hub	19
5	Hub port bypass	20
6	Arbitrated loop (FC-AL)	21
7	Loop after arbitration won and ports opened	21
8	First implementation on Symmetrix	22
9	Expanded connectivity	23
10	Dual loop disk drives	23
11	Highly available dual port disk solution	24
12	Loop addressing	25
13	AL_PA priority	26
14	AL_PA to loop ID chart	27
15	Primitive signals	28
16	Primitive sequences	29
17	Loop Port State Machine (LPSM)	31
18	Initialization procedure	34
19	LISM frame format	35
20	LIFA / LIPA / LIHA / LISA frames	37
21	LIRP/LILP frame format	40
22	FLOGI and Accept	41
23	PLOGI and Accept	42
24	PRLI and Accept	42
25	Arbitration Step 1	44
26	Arbitration Step 2	45
27	Arbitration Step 3	45
28	Arbitration Step 4	46
29	Arbitration Step 5	46
30	Arbitration Step 6	47

31	Access fairness window	48
32	Opening a loop circuit	49
33	Open primitive signals	50
34	Close loop: Step 1	51
35	Close loop: Step 2	51
36	Close loop: Step 3	52
37	Alternate credit	54
38	Single QuickLoop configuration	60
39	Dual QuickLoop configuration	61
40	Fabric/QuickLoop mixed-mode configuration	61
41	Simple storage mode configuration	66
42	Examples of tape pools	85
43	FC-AL high-availability distance topology example	92
44	FC-AL capacity expansion topology example	93
45	FC-AL high-availability capacity expansion topology example	93
46	FC-AL consolidation topology example	94
47	FC-AL high-availability consolidation topology example	94
48	FC-AL combined hub topologies example	95
49	Single-ID mode (virtualization disabled)	105
50	Multi-ID mode (virtualization enabled)	106
51	DataFort port mapping (storage side virtualization enabled)	107
52	DataFort LUN mapping (storage side virtualization enabled)	108
53	Recommended configuration of Decru DataFort with EMC storage ...	111
54	Topology without DataFort example	115
55	Virtualization enabled with port mapping	116
56	Recommended CryptoStor 2002 configuration for EMC storage products.....	120
57	CryptoStor Tape 700 deployment example	122

This document provides information on legacy SAN technologies.

E-Lab would like to thank all the contributors to this document, including EMC engineers, EMC field personnel, and partners. Your contributions are invaluable.

As part of an effort to improve and enhance the performance and capabilities of its product lines, EMC periodically releases revisions of its hardware and software. Therefore, some functions described in this document may not be supported by all versions 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.

Audience

This guide is intended for EMC field personnel, including technology consultants, and for the storage architect, administrator, and operator involved in acquiring, managing, operating, or designing a networked storage environment that contains EMC and host devices.

EMC Support Matrix and E-Lab Interoperability Navigator

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.

The *EMC Support Matrix* links within this document will take you to Powerlink where you are asked to log in to the E-Lab Interoperability Navigator. Instructions on how to best use the ELN (tutorial, queries, wizards) are provided below this **Log in** window. If you are unfamiliar with finding information on this site, please read these instructions before proceeding any further.

Under the **PDFs and Guides** tab resides a collection of printable resources for reference or download. All of the matrices, including the ESM (which does not include most software), are subsets of the E-Lab Interoperability Navigator database. Included under this tab are:

- ◆ The *EMC Support Matrix*, a complete guide to interoperable, and supportable, configurations.
- ◆ Subset matrices for specific storage families, server families, operating systems or software product.
- ◆ Host connectivity guides for complete, authoritative information on how to configure hosts effectively for various storage environments.

Under the **PDFs and Guides** tab, consult the *Internet Protocol* pdf under the "Miscellaneous" heading for EMC's policies and requirements for the *EMC Support Matrix*.

Related documentation

Related documents include:

- ◆ The former *EMC Networked Storage Topology Guide* has been divided into several TechBooks and reference manuals. The following documents, including this one, are available through the E-Lab Interoperability Navigator, **Topology Resource Center** tab, at <http://elabnavigator.EMC.com>.

These documents are also available at the following location:

<http://www.emc.com/products/interoperability/topology-resource-center.htm>

- *Backup and Recovery in a SAN TechBook*
- *Building Secure SANs TechBook*
- *Extended Distance Technologies TechBook*
- *Fibre Channel over Ethernet (FCoE): Data Center Bridging (DCB) Concepts and Protocols TechBook*
- *Fibre Channel SAN Topologies TechBook*
- *iSCSI SAN Topologies TechBook*
- *Networked Storage Concepts and Protocols TechBook*
- *Networking for Storage Virtualization and RecoverPoint TechBook*
- *WAN Optimization Controller Technologies TechBook*
- *EMC Connectrix SAN Products Data Reference Manual*
- *Non-EMC SAN Products Data Reference Manual*

- ◆ *EMC Support Matrix*, available through E-Lab Interoperability Navigator at <http://elabnavigator.EMC.com> > **PDFs and Guides**
- ◆ RSA security solutions documentation, which can be found at <http://RSA.com> > **Content Library**

All of the following documentation and release notes can be found at <http://Powerlink.EMC.com>. From the toolbar, select **Support > Technical Documentation and Advisories**, then choose the appropriate Hardware/Platforms, Software, or Host Connectivity/HBAs documentation links.

Hardware documents and release notes include those on:

- ◆ Connectrix B series
- ◆ Connectrix M series
- ◆ Connectrix MDS (release notes only)
- ◆ CLARiON
- ◆ Celerra
- ◆ Symmetrix

Software documents include those on:

- ◆ EMC Ionix ControlCenter
- ◆ RecoverPoint
- ◆ Invista
- ◆ TimeFinder
- ◆ PowerPath

The following E-Lab documentation is also available:

- ◆ Host Connectivity Guides
- ◆ HBA Guides

For Cisco and Brocade documentation, refer to the vendor's website.

- ◆ <http://cisco.com>
- ◆ <http://brocade.com>

Conventions used in this document



EMC uses the following conventions for special notices:

CAUTION

CAUTION, used with the safety alert symbol, indicates a hazardous situation which, if not avoided, could result in minor or moderate injury.



IMPORTANT

An important notice contains information essential to software or hardware operation.

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

Typographical conventions

EMC uses the following type style conventions in this document.

Normal	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> 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, functions, utilities URLs, pathnames, filenames, directory names, computer names, filenames, links, groups, service keys, file systems, notifications
Bold	Used in running (nonprocedural) text for: <ul style="list-style-type: none"> Names of commands, daemons, options, programs, processes, services, applications, utilities, kernels, notifications, system calls, man pages Used in procedures for: <ul style="list-style-type: none"> Names of interface elements (such as names of windows, dialog boxes, buttons, fields, and menus) What user specifically selects, clicks, presses, or types
<i>Italic</i>	Used in all text (including procedures) for: <ul style="list-style-type: none"> Full titles of publications referenced in text Emphasis (for example a new term) Variables
<code>Courier</code>	Used for: <ul style="list-style-type: none"> System output, such as an error message or script URLs, complete paths, filenames, prompts, and syntax when shown outside of running text
<code>Courier bold</code>	Used for: <ul style="list-style-type: none"> Specific user input (such as commands)
<i><code>Courier italic</code></i>	Used in procedures for: <ul style="list-style-type: none"> Variables on command line User input variables
< >	Angle brackets enclose parameter or variable values supplied by the user

[]	Square brackets enclose optional values
	Vertical bar indicates alternate selections - the bar means “or”
{ }	Braces indicate content that you must specify (that is, x or y or z)
...	Ellipses indicate 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 Powerlink and choose **Support**. On the Support page, you will see several options, including one for making a service request. Note that to open a service request, you must have a valid support agreement. Please contact your EMC sales representative for details about obtaining a valid support agreement or with questions about your account.

We'd like to hear from you!

Your feedback on our TechBooks is important to us! We want our books to be as helpful and relevant as possible, so please feel free to send us your comments, opinions and thoughts on this or any other TechBook:

TechBooks@emc.com

Fibre Channel Arbitrated Loop (FC-AL)

This chapter contains information on Fibre Channel arbitrated loop (FC-AL).

- ◆ FC-AL..... 16
- ◆ Symmetrix and Fibre Channel connectivity..... 22
- ◆ Arbitrated loop addressing..... 25
- ◆ Primitive signals and sequences 28
- ◆ Loop initialization..... 33
- ◆ Arbitration process 43
- ◆ Alternate Buffer-to-Buffer Credit (BB_Credit) 53

FC-AL

This section contains the following information:

- ◆ “Overview” on page 16
- ◆ “Loop construction” on page 18
- ◆ “Loops with hubs” on page 19
- ◆ “Arbitration” on page 20

Overview

When Fibre Channel was first introduced, it was a new technology and everything was expensive. Switches, hubs, and node transceivers proved to be costly. Arbitrated Loop topology lies between point-to-point and switched fabric in that it provides more connectivity than point-to-point with up to 126 NL_Ports in a loop, but less than switched fabric which has the ability in theory to support up to 16 million ports. It was a cost-effective way of connecting a limited number of ports in a loop single network.

Fibre Channel arbitrated loop (FC-AL) is a *daisy-chain* connecting up to 126 devices in a loop configuration over attachment points called *L_Ports* (loop ports). FC-AL is a low-cost connectivity solution because it does not require switches. FC-AL is a good choice for small to medium-sized configurations, and provides a growth path by allowing connection of a loop to a switched fabric.

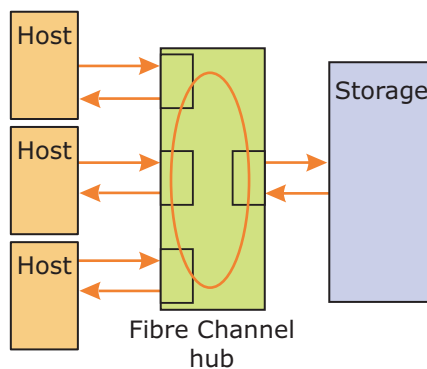


Figure 1 Arbitrated loop example

Efficiency and connectivity is enhanced by incorporating one or more hubs into the loop. Routing traffic through a hub on each leg of a loop eliminates the loss of the entire loop, as happens in a hubless loop. (“Loops with hubs” on page 19 provides more information.)

The arbitrated loop topology promoted the introduction of Fibre Channel by removing the cost of a fabric switch and, depending on the number of nodes in the loop, the amount of transceivers could also be reduced.

Arbitrated loop provides more connectivity than point-to-point in that it can support 126 NL_Ports and 1 FL_Port on a single loop and is a middle ground between point-to-point and switched fabric. In arbitrated loop, the transmit output of one port is connected to the receive input of the next and these connections are made between all the nodes until a closed loop is formed (refer to Figure 2 on page 17). This type of configuration is usually made using a Fibre Channel hub which eliminates the need to form the logical loop using cabling.

In arbitrated loop each port sees all messages on the loop and ignores or passes those messages which are not addressed to that particular port.




Figure 2 Arbitrated loop

Loop construction

An arbitrated loop is constructed by connecting nodes together in a single loop. Loops can be constructed by physically connecting each node in the loop or through the use of a Fibre Channel hub. The transmit of one port is attached to the receive of the next. This is continued until the loop has been formed. A loop can contain 126 N_L Ports and one F_L Port which is used for connection to a Fabric environment. Information that is passed around the loop is repeated by each port and passed on if it is not the required destination.

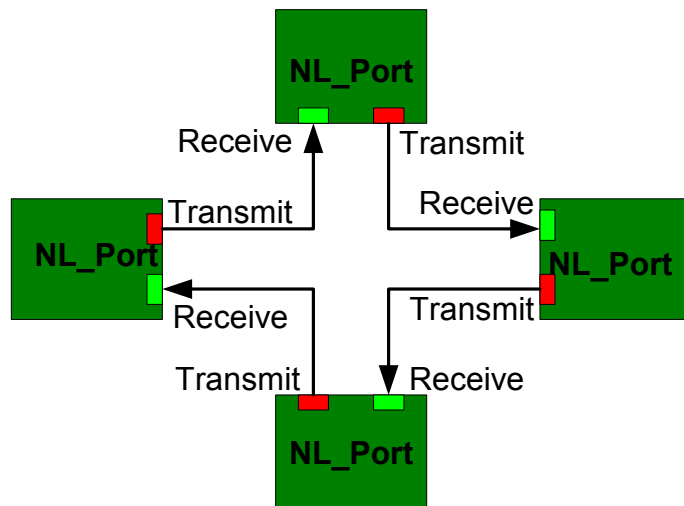


Figure 3 Arbitrated loop with 4 nodes

Figure 3 shows an example of how an arbitrated loop might be constructed. In this loop, there are four nodes and four transceivers, keeping the Fibre Channel circuitry to a minimum. In this configuration, the bandwidth is shared between all nodes on the loop and if any new nodes were added it would be further reduced. In this configuration if any of the ports failed then the complete loop would be lost. With the absence of a hard failure, this loop could be thrown into turmoil if any intermittent type failure should arise. Any *blip* whatsoever would force the whole loop into a re-initialization state and halt all I/O operations until the loop is back to a stable state. This is not desirable in a highly-available storage subsystem which is why today most fibre connectivity to a Symmetrix is through a fabric switch (refer to “Symmetrix and Fibre Channel connectivity” on

page 22). However, with the emergence of dual port fibre devices, the arbitrated loop is a viable option, especially as a back-end solution within a storage array.

Arbitrated loop performance is dependant on a number of factors. Some of these are obvious, such as the number of nodes on the loop. Clearly, in a shared bandwidth medium, the population fighting for that medium will have a direct impact on performance. However, this is not the only way the number of nodes can affect performance.

With the introduction of more nodes, the roundtrip time of the loop is also increased. This is because each frame or sequence may have additional ports to pass through to reach the final destination, and with each additional port adding some latency, overall throughput can be affected. Likewise, the time it takes to win arbitration can increase.

Loops with hubs

Another way of constructing a loop is to use a hub. As the cost of transceivers and hubs have dropped dramatically from when they were first introduced, the benefits obtained with this configuration can outweigh the extra expense. As shown in Figure 4, the loop is constructed within the hub.

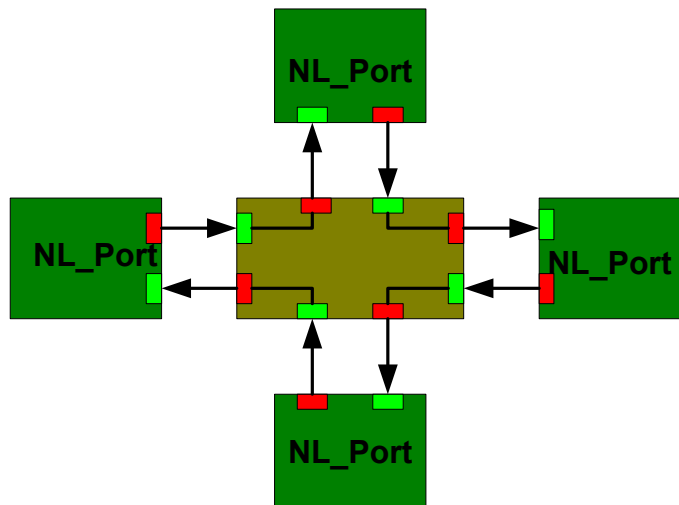


Figure 4 Arbitrated loop with a hub

The first improvement is the ease of cabling. In [Figure 4](#), each transmitter had to be routed to the receiver of the next node and so on. As shown in [Figure 4](#), all that is needed to construct the loop is to plug each cable directly to the hub. This example shows that for the same number of nodes in the loop, the number of transceivers will double, adding to the cost.

The hub brings another useful feature to the table. A failure would cause the whole loop to collapse. However, the hub can bypass a faulty port and allow the remaining ports to operate as normal, as illustrated in [Figure 5](#).

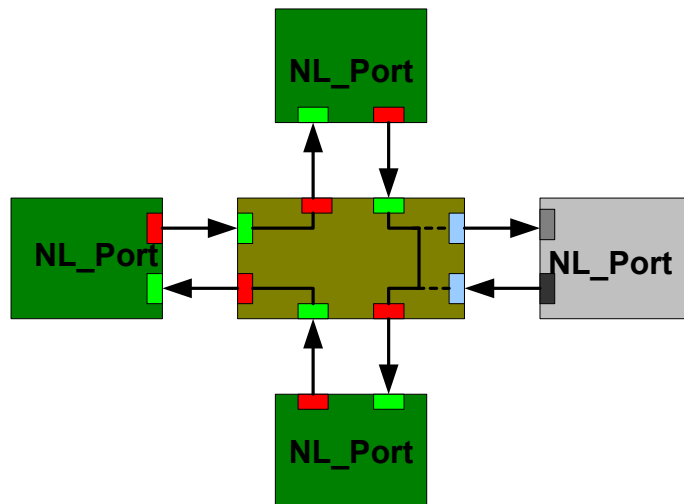


Figure 5 Hub port bypass

Arbitration

The process through which a fibre port gains sole control of the loop is called *arbitration*. Once a port has control of the loop, and it has opened another port, it is free to transmit frames to that port. [Figure 6 on page 21](#) shows a loop which is available; in other words devices are free to arbitrate to gain access to this loop.




Figure 6 Arbitrated loop (FC-AL)

Arbitration prevents multiple ports from sending frames at the same time on the loop. After a device has won arbitration and opened another device then, in effect, a point-to-point connection has been established. All remaining ports cannot participate in any way other than passing along whatever comes their way until the loop has been closed and is once again available for arbitration to all.

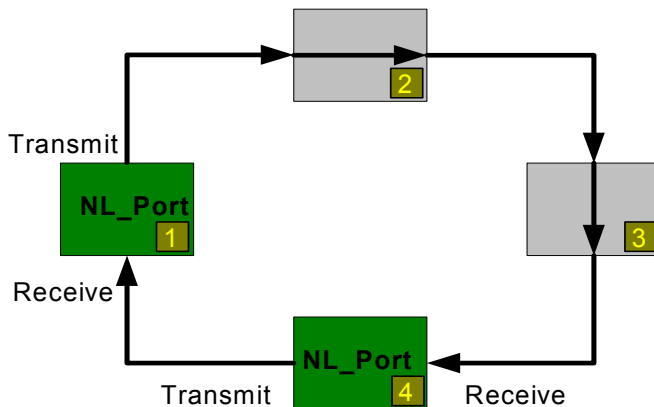


Figure 7 Loop after arbitration won and ports opened

Figure 7 illustrates the state of the loop after a successful arbitration by port one or four. The arbitration winner opened the other port and now ports two and three have been logically removed from the loop forming a virtual point-to-point connection between port one and four.

Symmetrix and Fibre Channel connectivity

This section contains the following information:

- ◆ “Overview” on page 22
- ◆ “Dual port devices” on page 23

Overview

EMC first introduced arbitrated loop in a direct connect from HBA to Symmetrix FA configuration (2 Node Arbitrated Loop). (Refer to [Figure 8](#)). This was the beginning of Fibre Channel connectivity to the Symmetrix and the configuration was limited to a minimum until all of the teething problems associated with a new technology had been identified and corrected.




Figure 8 First implementation on Symmetrix

Very quickly, hubs were introduced to the configuration and the connectivity was expanded ([Figure 9 on page 23](#)). The hubs provided resiliency in that they provided a bypass circuit on each port which is used to bypass a particular port if there is not a valid signal present, allowing the rest of the loop to continue to operate. Basically, they ensure that a bad port or a broken fibre will not take down the entire loop.

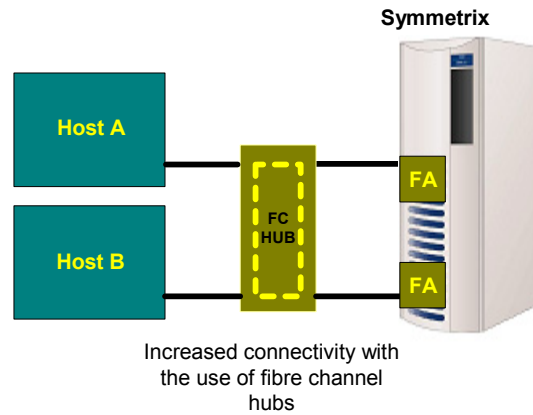


Figure 9 Expanded connectivity

Dual port devices

In dual loop drive systems each disk drive has two ports and thus the drive can be connected on two separate loops. As shown in [Figure 10](#), if a failure occurs on one loop, for example Initiator 1 fails or a port fails on one of the disks, the disks will remain accessible through the second initiator. This could be compared to PowerPath at the disk level.




Figure 10 Dual loop disk drives

This is not the only solution available with dual port disks. Hubs could also be included which would provide further connectivity and resiliency options (see [Figure 11](#)).

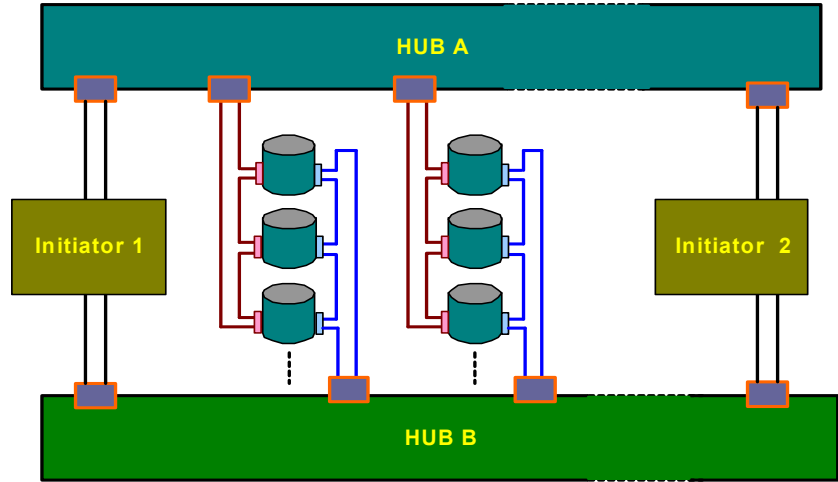


Figure 11 Highly available dual port disk solution

Arbitrated loop addressing

This section contains the following information:

- ◆ “Overview” on page 25
- ◆ “Loop ID” on page 26

Overview

Fibre Channel specifies a three-byte field for the address used in routing frames. In arbitrated loop, only one of these three bytes (least significant 8 bits) is used for the address which is known as the Arbitrated Loop Physical Address (AL_PA). This address is used in the Source ID (S_ID) and Destination ID (D_ID) of frames transmitted in the loop.

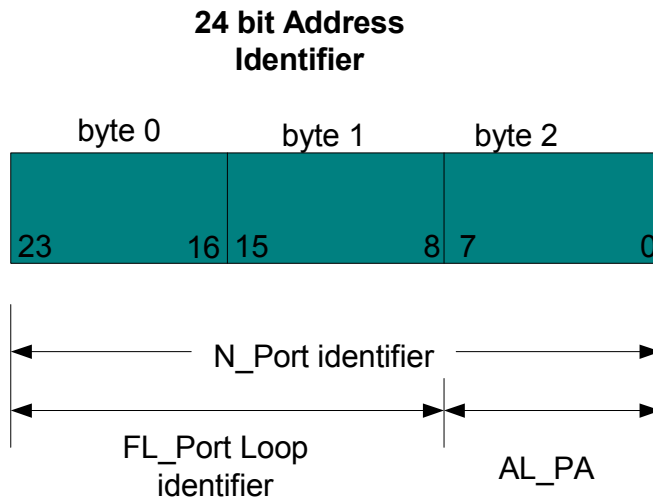


Figure 12 Loop addressing

Figure 12 shows the full 24-bit address defined by the Fibre Channel standard. Eight (8) bits used by the AL_PA. Bits 8 to 23 are used for the FL_Port identifier and the full 24 bits are used by an N_Port in a fabric switch environment.

The AL_PA values used are limited to characters that result in neutral disparity after encoding. AL_PA 00 is reserved for FL_Port and the remaining 126 AL_PA are distributed irregularly between 01 and EF. If two ports arbitrate to get access to the loop at the same time then

only one port can be given access. This is decided by assigning priorities to the AL_PA addresses with 00 having the highest and 01 through to EF having decreasing priority.

Figure 13 shows AL_PA priorities.

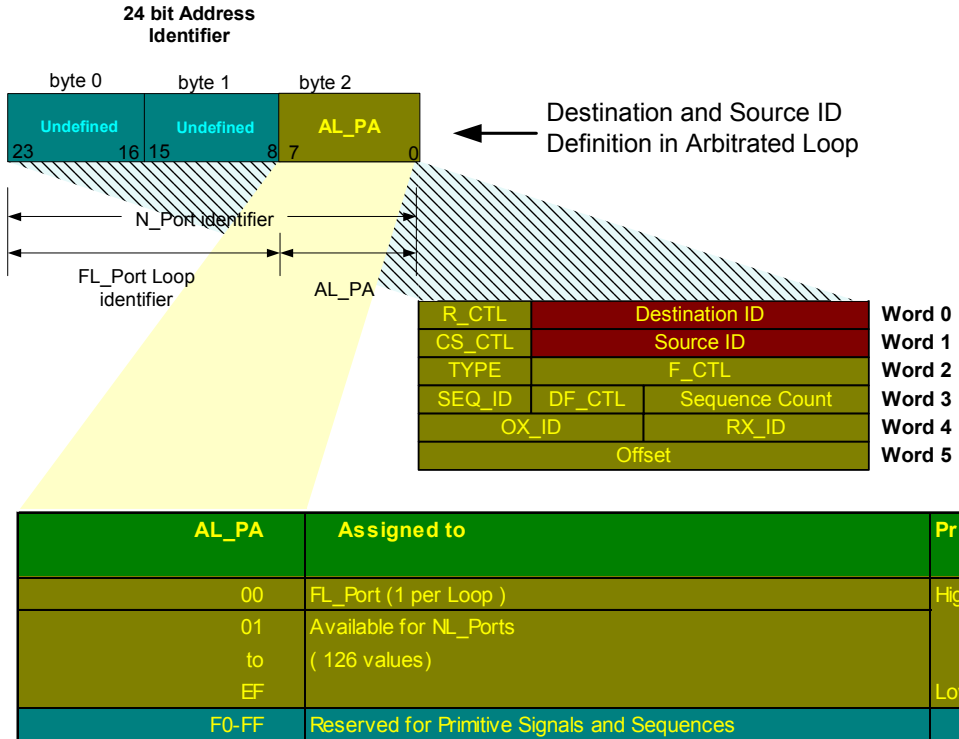


Figure 13 AL_PA priority

Loop ID

The usable AL_PA values are made up of the neutral disparity characters in the range 0x00 to 0xFF and are not distributed regularly. The used AL_PA values range from 0x00 to 0xEF as is shown in Figure 14 on page 27. When you configure a Symmetrix FA to have a certain AL_PA you do not set an AL_PA value directly but instead specify a Loop ID value in the bin file and this Loop ID then corresponds to a certain AL_PA value. The Loop ID values are sequential between 0x00 and 0x7E and thus can be easier to manage. The correlation of Loop ID to AL_PA is shown in Figure 14.

AL_PA	Loop ID		AL_PA	Loop ID	
	HEX	Decimal		HEX	Decimal
EF	0	0	72	40	64
E8	1	1	71	41	65
E4	2	2	6E	42	66
E2	3	3	6D	43	67
E1	4	4	6C	44	68
E0	5	5	6B	45	69
DC	6	6	6A	46	70
DA	7	7	69	47	71
D9	8	8	67	48	72
D6	9	9	66	49	73
D5	0A	10	65	4A	74
D4	0B	11	63	4B	75
D3	0C	12	5C	4C	76
D2	0D	13	5A	4D	77
D1	0E	14	59	4E	78
CE	0F	15	56	4F	79
CD	10	16	55	50	80
CC	11	17	54	51	81
CB	12	18	53	52	82
Ca	13	19	52	53	83
C9	14	20	51	54	84
C7	15	21	4E	55	85
C6	16	22	4D	56	86
C5	17	23	4C	57	87
C3	18	24	4B	58	88
BC	19	25	4A	59	89
BA	1A	26	49	5A	90
B9	1B	27	47	5B	91
B6	1C	28	46	5C	92
B5	1D	29	45	5D	93
B4	1E	30	43	5E	94
B3	1F	31	3C	5F	95
B2	20	32	3A	60	96
B1	21	33	39	61	97
A E	22	34	36	62	98
AD	23	35	35	63	99
AC	24	36	34	64	100
AB	25	37	33	65	101
AA	26	38	32	66	102
A9	27	39	31	67	103
A7	28	40	2E	68	104
A6	29	41	2D	69	105
A5	2A	42	2C	6A	106
A3	2B	43	2B	6B	107
9F	2C	44	2A	6C	108
9E	2D	45	29	6D	109
9D	2E	46	27	6E	110
9B	2F	47	26	6F	111
98	30	48	25	70	112
97	31	49	23	71	113
90	32	50	1F	72	114
8F	33	51	1E	73	115
88	34	52	1D	74	116
84	35	53	1B	75	117
82	36	54	18	76	118
81	37	55	17	77	119
80	37	56	10	78	120
7C	39	57	0F	79	121
7A	3A	58	8	7A	122
79	3B	59	4	7B	123
76	3C	60	2	7C	124
75	3D	61	1	7D	125
74	3E	62	0	7E	126
73	3F	63			

Figure 14 AL_PA to loop ID chart

Primitive signals and sequences

Arbitrated loop has several *ordered sets* used in loop arbitration and opening and closing of loop circuits (refer to “[Opening and Closing the loop](#)” on page 49). An ordered set is a group of four transmission characters, the first being the K28. Five special characters and the remaining three (data) characters define the meaning of the ordered set. These ordered sets can exist either at the start or end of the frame (in the case of frame delimiters), or can exist on their own (in the case of primitive signals and sequences). For more information on ordered sets, refer to the “Ordered sets” section in the *Networked Storage Concepts and Protocols TechBook*, available through the E-Lab Interoperability Navigator, **Topology Resource Center** tab, at <http://elabnavigator.EMC.com>.

Primitive signals are normally used to indicate events or actions at the sending port. A table of these primitive signals is shown in [Figure 15](#). For example, the ARB(x) primitive signal is used by a port in the arbitrating state to indicate that it is arbitrating for access to the loop. The x indicates the AL_PA assigned to that specific port.

Primitive Signal	Abbr	Ordered Set
Arbitrate	ARB(x)	K28.5 D20.4 AL_PA AL_PA
Arbitrate Fairness	ARB(F0)	K28.5 D20.4 D16.7 D16.7
Arbitrate (No AL_PA)	ARB(F7)	K28.5 D20.4 D23.7 D23.7
Clock Synchronization X	SYN x	K28.5 D31.3 CS_x CS_x
Clock Synchronization Y	SYN y	K28.5 D31.5 CS_y CS_y
Clock Synchronization Z	SYN z	K28.5 D31.6 CS_z CS_z
Close	CLS	K28.5 D05.4 D21.5 D21.5
Dynamic Half Duplex	DHD	K28.5 D10.4 D21.5 D21.5
Idle	IDLE	K28.5 D21.4 D21.5 D21.5
Mark	MRK(x)	K28.5 D31.2 MK_TP AL_PS
Open Full-Duplex (Point-to-Point)	OPN(yx)	K28.5 D17.4 AL_PD AL_PS
Open Half-Duplex (Point-to-Point)	OPN(yy)	K28.5 D17.4 AL_PD AL_PD
Open Broadcast Replicate	OPN(fr)	K28.5 D17.4 D31.7 D31.7
Open Selective Replicate	OPN(yr)	K28.5 D21.4 AL_PD D31.7
Receiver Ready	R_RDY	K28.5 D21.4 D10.2 D10.2

Figure 15 Primitive signals

Primitive sequences are used to indicate states or conditions and are normally transmitted continuously until something causes the state to change. A minimum of three consecutive occurrences of the same ordered set is required before a primitive sequence is recognized and action taken. Figure 16 shows Arbitrated Loop primitive sequences used in link initialization and loop port bypass. Switched fabric uses a different method of link initialization than arbitrated loop, but it does still use primitive sequences.

Primitive Sequence	Abbr	Ordered set
Loop Initialization - F7,F7	LIP	K28.5 D21.0 D23.7 D23.7
Loop Initialization - F8,F7	LIP	K28.5 D21.0 D24.7 D23.7
Loop Initialization - F7,x	LIP	K28.5 D21.0 D23.7 AL_PS
Loop Initialisation - F8,x	LIP	K28.5 D21.0 D24.7 AL_PS
Loop Initialization - reset	LIPyx	K28.5 D21.0 AL_PD AL_PS
Loop Port Enable	LPEyx	K28.5 D5.0 AL_PD AL_PS
Loop Port Enable All	LPEfx	K28.5 D5.0 D31.7 AL_PS
Loop Port Bypass	LPByx	K28.5 D9.0 AL_PD AL_PS

Figure 16 Primitive sequences

An AL_PA identifies either a source or destination port in arbitrated loop. In some cases it is necessary to identify whether it is the source or destination AL_PA that is being referred to. When it is necessary to identify a destination port, the term **AL_PD** is used. In the case of the source port, the term **AL_PS** is used. Figure 16 shows different types of LIP (Loop Initialization Primitive) sequence used in loop initialization. Each is discussed below.

Loop Initialization LIP(F7, F7). A port transmitting LIP(F7, F7) indicates that the port in the initializing state is requesting loop initialization but does not have a valid AL_PA. This mainly occurs when a device is hot-plugged into a Loop or when a port that was non-participating wants to become participating and requires an AL_PA to do so.

Loop Failure LIP(F8, F7). A port transmitting LIP(F8, F7) indicates that the port in the initializing state is requesting a loop initialization due to a loop failure. The port does not have an AL_PA and uses F7. This could occur if a non-participating loop port without an AL_PA detects a loop failure or a node in the process of getting an AL_PA detects a loop failure.

Loop Initialization LIP(F7, AL_PS). This LIP indicates that the loop port identified in the AL_PS value is requesting loop initialization. This can be used if the port detects a performance degradation, arbitration wait time-out, or for another unspecified reason.

Loop Failure LIP(F8, AL_PS). This LIP indicates that the loop port identified in the AL_PS value has detected a loop failure. This may occur when a loop interconnection has failed, a loop port has failed, a loop port has been powered off or removed from the loop when no bypass circuit is present, or if the bypass circuit fails.

Selective Reset LIP(AL_PD, AL_PS). The selective reset LIP is used to perform a vendor specific reset at the loop port specified in the AL_PD value. The AL_PS value indicates the port that originated the request. This LIP could be used during error recovery to reset a port that is in a hung state.

Along with the loop initialization primitive sequences, there are a number of sequences to set and reset the LP_Bypass variable in the Loop Port State Machine (LPSM). With this variable set, the LPSM retransmits frames and does not attempt to arbitrate or participate in the loop. The state of this variable is also used to control an optional port bypass circuit to electrically bypass the loop port.

Loop Port Enable LPE(yx), LPE(fx). These primitive sequences cause either a designated port (yx) or all ports (fx) to reset the LP_Bypass variable and deactivate a control line to an external port bypass circuit if present.

Loop Port Bypass LPB(yx), LPB(fx). These primitive sequences, when received, cause the designated port (yx) or all the ports (fx) to set the LP_Bypass variable and optionally activate a control line to an external port bypass circuit if present.

Loop Port State Machine (LPSM)

A port on an arbitrated loop is required to behave in a certain manner in order to operate correctly on the Loop. The LPSM defines this behavior (refer to [Figure 17 on page 31](#)). The particular states that can occur on a loop include: initialization, arbitration, open circuit, close circuit, and implement fairness. These various states are implemented using specific ordered sets which are processed by the LPSM of the ports on the loop.

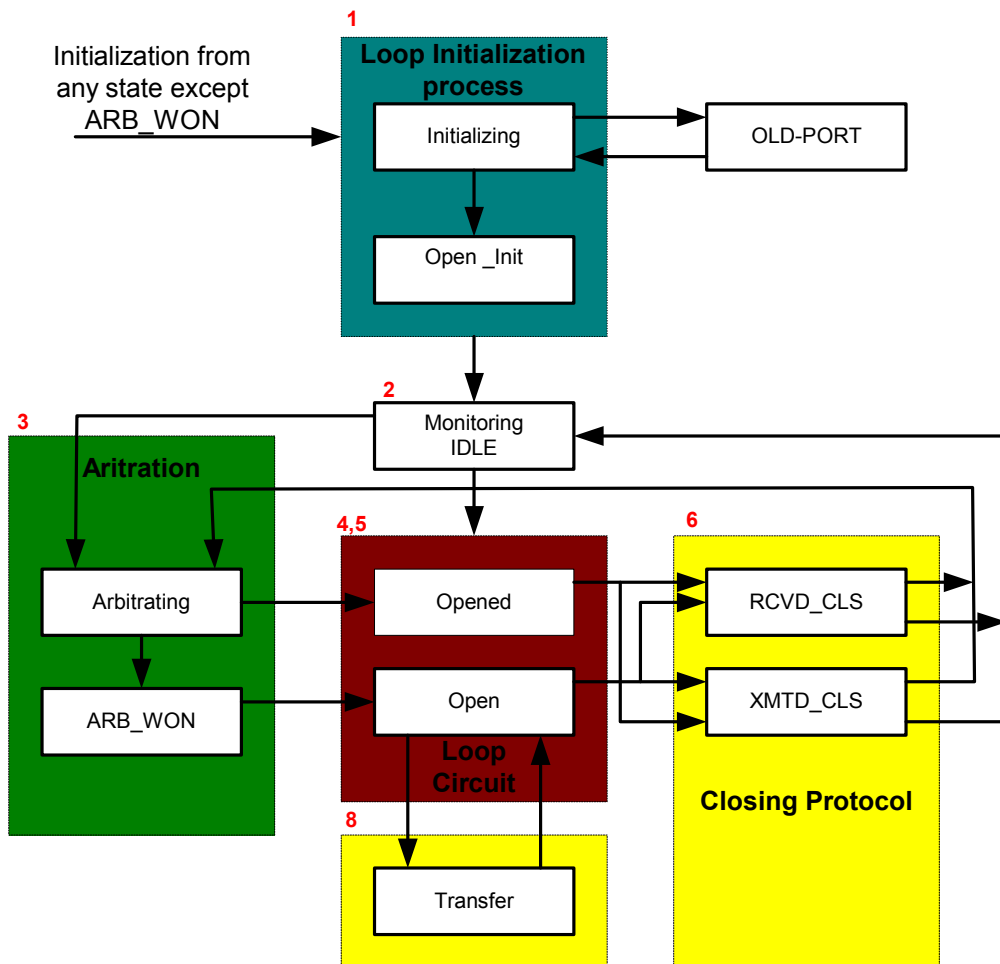


Figure 17 Loop Port State Machine (LPSM)

The basic operations of a Loop Port State Machine are:

- ◆ When a port is first attached to the loop it enters the initializing state and starts loop initialization where each port transmits a LIP continuously and monitors for a LIP returning. On receipt of a LIP the ports go into a Open-Init state where most of the Loop initialization steps take place. If the loop is not operational the port may go to into the Old_Port state disabling Arbitrated Loop functionality and begin operating as an N_Port.

- ◆ If the ports have no work to do they are in the monitoring state transmitting and receiving Idles.
- ◆ When the loop port requires access to other ports in the loop it enters the arbitrating state and begins arbitrating for access. When the port wins arbitration it enters the ARB-WON state.
- ◆ When the port needs to access another port on the loop it sends an OPN to select the destination port. The source port goes into the OPEN state and the destination port goes into the OPENED state.
- ◆ Once both ports are in these states transfer of frame can begin.
- ◆ When either of the ports has completed its transfer and wants to close the connection it sends a close (CLS). The port that sends the CLS enters the transmitted close state and the when the other port receives the CLS it enters the received close state. It transmits its remaining frames, if necessary, as long as it has available credit and then transmits a CLS back to the originator. The two ports are now logically disconnected from the loop and enter the monitoring state again.
- ◆ If a loop port is in the monitoring state or arbitrating for the loop and receives an OPEN from another port it enters the OPENED state.
- ◆ There is another state that a port can enter called the *transfer state*. If a port has frames to send to multiple ports and is finished communicating with the first of these ports it can send a CLS, move to the transfer state, and once it receives the CLS from the first port it can open the second port without going through an arbitration cycle.

Loop initialization

Before discussing all the steps in detail that take place during a loop initialization, it is important to understand why this step must be done and exactly what is accomplished during this step.

This section contains the following information:

- ◆ [“Overview” on page 33](#)
- ◆ [“Loop initialization steps” on page 36](#)
- ◆ [“Login process” on page 40](#)

Overview

The loop initialization step performs a number of functions in arbitrated loop including the assignment of addresses (AL_PAs) to loop ports, notification that the configuration may have changed, and notification of a loop failure. Events that can cause a loop initialization include: a) if a port was powered on it may need to acquire an AL_PA and notify other ports on the loop that the configuration has changed; or b) a port detects a physical connection problem and begins the loop initialization process to notify other ports and check if the loop is still operational. The main steps involved in loop initialization are listed in [Figure 18](#).

Steps

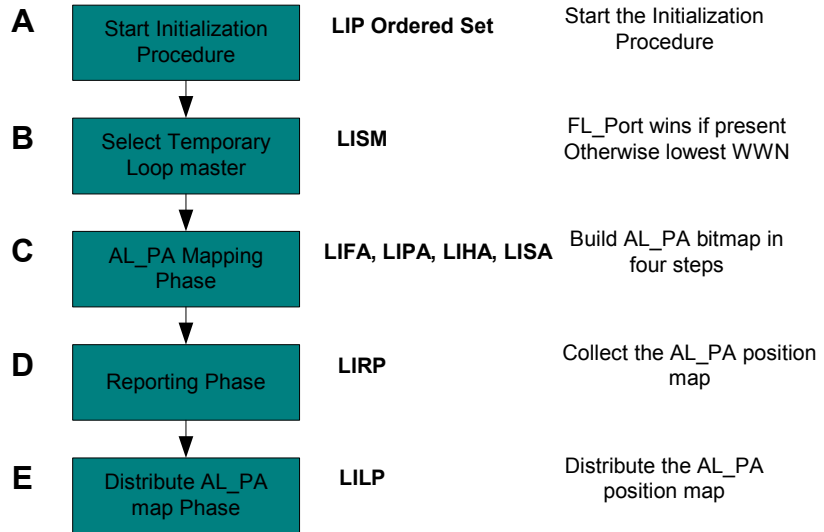


Figure 18 Initialization procedure

Each node on a loop is capable of initiating loop initialization by transmitting one of the LIP sequences shown in [Figure 16 on page 29](#). When a node transmits this LIP sequence all the other ports on the loop recognize this and enter the **OPEN_INIT** state and retransmit the LIP along the loop. This occurs until the LIP arrives back at the port that is initiating the initialization and it too enters the 'OPEN_INIT' state.

With all ports on the loop now in the OPEN_INIT state, the next step is a process to select a port to become the temporary loop master. This step is initiated by each node on the loop that has entered 'OPEN_INIT' state continuously sending out LISM (Loop Initialization Select Master) frames. LISM frames have a certain format, detailed [Figure 19 on page 35](#), and serve the purpose of allowing each node on the loop a chance at becoming the temporary loop master.




Figure 19 LISM frame format

The structure of the LISM frame is shown in [Figure 19](#) with the S_ID and D_ID fields in the frame header showing a value of '0000xx'. The least significant byte is a value of 00 if the port is an FL_Port, or EF if the port is an NL_Port which identifies the port type. The payload of the LISM frame contains the 8 byte World Wide Name of the port.

If there is a single FL_Port on the loop it becomes the loop master. If there are more than one FL_Port then the one with the lowest port worldwide number becomes the master. The reason the FL_Port becomes the master is based on the fact that the FL_Port is part of a fabric and therefore has more knowledge of the configuration. If there are no FL_Ports present then the NL_Port with the lowest port worldwide number is selected as temporary loop master. Normally, it is an HBA that becomes the loop master as its worldwide number is lower than that of a Symmetrix FA.

The process involved at the LISM stage involves each port transmitting LISM frames and each port checking the LISM frames it is receiving for the port type field which is denoted in the least significant byte of the S_ID and D_ID fields. If an NL_Port receives a LISM frame from an FL_Port it stops transmitting its own LISM frames and begins retransmitting the LISM of the FL_Port. If an FL_Port receives a LISM frame from an NL_Port it discards the received frame and begins transmitting its own frame. If the port type in the S_ID and D_ID fields of the frame header is the same as that of the receiving port, the port worldwide number in the payload of the LISM frame is compared with that of the receiving port. If the port worldwide number in the received LISM frame is higher than that of the receiving port, the port discards the frame and the port continues to transmit its own LISM. If the port worldwide number in the received LISM frame is not higher, it stops transmitting its own LISM frame and begins transmitting the received frame. Eventually one of

the loop ports will receive back around the loop its own LISM frame and when this happens this port becomes the temporary loop master. This port then begins transmitting the ARB(F0) primitive signal to inform the other ports that the LISM procedure has been completed and a loop master selected.

The next step in the process is *AL_PA assignment*. This process involves the assignment of Arbitrated Loop Physical Addresses (AL_PA) to each port on the loop. A port's AL_PA is its Fibre Channel address on the loop and is used to identify it on the loop and is put in the S_ID of the frame header of all data frames that are transmitted by the port and is in the D_ID of all data frames that are to be received by the port. This is a multistep process where addresses are assigned using four distinct steps depending on the particular ports operation. These four steps are discussed beginning on [page 37](#) and all use the concept of populating an AL_PA bitmap with a value depending on whether that corresponding AL_PA is assigned or not.

In arbitrated loop there are 127 possible addresses on a loop (126 for NL_Ports and one for the FL_Port). To identify which AL_PA values have been taken in a loop a 128 bit (four word) map is used where each bit corresponds to a certain AL_PA. If a bit is set to 1 then the corresponding AL_PA is assigned and that address is in use. If the bit is 0 then the corresponding AL_PA is available to be acquired. Word 0 bit 31 of the AL_PA bit map is the login-required bit (L-bit) which is set by an FL_Port to indicate that the configuration has changed and that all ports are logged out.

Loop initialization steps

This section describes the loop initialization steps. [Figure 20 on page 37](#) shows LIFA / LIPA / LIHA / LISA frames.




Figure 20 LIFA / LIPA / LIHA / LISA frames

Step 1. The first step in the AL_PA assignment is the LIFA (Loop Initialization Fabric Address) which basically allows public ports that had previously been logged in with the fabric (FL_Ports) to reclaim the AL_PA they had been using prior to loop initialization. The temporary loop master first initializes the AL_PA bitmap to all '0's and builds the LIFA frame with this bitmap in payload words two through five and the loop initialization identifier in word 1, as is shown in Figure 20.

If the loop master had a fabric assigned address prior to the loop initialization it sets the bit for that AL_PA in the bitmap and then transmits this frame. If the next port on the loop had a fabric assigned AL_PA it also sets the corresponding bit in the AL_PA bitmap (if not already set) and retransmits the frame. If the next loop port did not previously have a fabric assigned AL_PA then it would leave the AL_PA map unchanged and retransmits the frame. This continues until every port on the loop processes the LIFA frame and is received back by the temporary loop master.

Step 2. The next step of the process is **LIPA** (Loop Initialization Previous Address) where private ports (NL_Ports) that had an AL_PA prior to loop initialization can reclaim the same AL_PA. This step is initiated by the temporary loop master by changing the loop initialization identifier in word 1 of the LIFA frame it received to LIPA identifier and then transmitting this around the loop. Each port starting with the loop master checks to see if it had a nonfabric assigned AL-PA prior to loop initialization. If it had, it first checks to see if that bit is set in the AL_PA bit map and if it is not set then sets it to 1 and reclaims the AL_PA. If the bit is set then the port would have to wait for a soft-assigned address. This continues until every port on the loop has processed the frame and it is received back by the loop master.

Step 3. The next step of the initialization process is for ports that did not have an AL-PA prior to loop initialization but do have a preferred AL_PA that is set by either jumper settings or some other configuration method. This step is known as the **LIHA** (Loop Initialization Hard Address)/ The Symmetrix FA would be an example where the AL_PA can be configured in the bin file by setting a corresponding Loop ID value. The temporary loop master changes the received LIPA frame loop initialization identifier in word 1 of the payload to LIHA which indicates this is now a LIHA frame which will be transmitted around the loop. Each port that did not have an AL_PA prior to link initialization but does have hard assigned AL_PA checks the bit corresponding to that AL_PA in the LIHA frame payload words two to five and then sets the bit and claims that AL_PA. If the bit is already set then the port would again need to wait for a soft-assigned address. This continues until every port has processed the LIHA frame and it is received back around the loop by the temporary loop master.

Step 4. The final step in acquiring an AL_PA is the **LISA** (Loop Initialization Soft Address) where the port may select the first available AL_PA in the bitmap. Once the loop master has received back the LIHA frame it changes the identifier in word 1 from LIHA to LISA and leaves the AL_PA map unchanged. It then transmits the frame and every port that does not have an AL_PA scans the bitmap in the LISA frame to find the first available AL_PA and then claims that soft-assigned address by setting that bit in the AL_PA bitmap. The AL_PA values are normally assigned starting at the most significant bit of the AL_PA map and proceeding to the least significant bit. This provides the most efficient ordering of AL_PA values around the loop. For the best performance the AL_PA values

should be arranged in descending priority in the direction of information flow around the loop. AL_PA 00 is the highest priority.

If after these steps a port has not acquired an address, for example if all the available AL_PA s have been taken or if a ports hard address is already in use, then the port will enter the non- participating mode.

If the port does not support loop positional mapping steps which follow AL_PA assignment it sets the third byte of the loop initialization identifier (Word 1) to 00.

A method assigning addresses during the LISA step is to have initiators acquire AL_PA values in the higher priority end of the AL_PA bitmap starting at the most significant bit and proceeding to the least significant bit and targets do the opposite. However, this does not lead to the most efficient ordering of the AL_PAs on a loop for target devices as they would get AL_PAs assigned starting with the lower priority.

A preferred method of assigning AL_PA values during the LISA process is to have a range of AL_PA values at the higher order end of the AL_PA bit map reserved for initiators and have targets assign soft addresses beginning at the end of this range. This would avoid the AL_PA ordering problem in that the initiators would get the higher priority AL_PA values in descending order and the targets would get the lower priority AL_PA values also in descending order around the loop.

There are two additional steps in the initialization stage that are needed to provide information on the positioning of AL_PAs in the loop for managing the configuration and problem analysis. These two steps are Loop Initialization Report Position (LIRP) and Loop Initialization Loop Position (LILP):

- ◆ The **LIRP** step involves building a map of the AL_PA values according to their position on the loop relative to the temporary loop master. The temporary loop master begins this process by building a LIRP frame with the structure shown in [Figure 21 on page 40](#).




Figure 21 LIRP/LILP frame format

The basic structure of the frame payload is a one word loop initialization identifier, a one byte offset and up to 127 AL_PA entries. The temporary loop master first of all initializes the structure by setting the offset to 01 and storing its AL_PA at offset 01 in the AL_PA map and puts FF in all the other remaining positions. This frame is then passed to the next port on the loop, the offset is incremented by 1 and the next port stores its AL_PA at that location in the map. This process continues until the LILP frame comes back around the loop to the temporary loop master.

- ◆ When the loop master receives the LIRP frame back it changes the identifier in the payload of the frame to a **LILP** and retransmits the entire frame so each port can have a copy of the AL_PA positional map. Once this frame comes back around the loop the loop initialization is complete.

Login process

Once the loop is initialized each port has acquired an address but each initiator (HBA in server) does not know what target devices are on the loop. For each HBA to discover what targets are on the loop it needs to perform some extra steps. These steps provide ports with a means of exchanging information about each other that is used to control any communication that is initiated by these ports. Fibre

Channel provides three different types of login that can occur between ports by the use of the following extended link services:

- ◆ Fabric Login (FLOGI)
- ◆ N_Port Login (PLOGI)
- ◆ Process Login (PRLI)

The **Fabric login** process is used in a switched fabric environment to allow an N_Port establish a session with the fabric. During this step both the N_Port and the fabric exchange parameters with each other and identify themselves to each other. This step also assigns an address to the attached N_Port.




Figure 22 FLOGI and Accept

Figure 22 shows the exchange of information that takes place with the N_Port sending the initial FLOGI frame (Command code 04) with all its information and the Fabric returning the Accept (Command code 02) with its associated parameters.

In arbitrated loop the Fabric login does not take place and the ports exchange service parameters with **N_Port login**, which is also known as PLOGI. Service parameters are basically information regarding the FC-2 capabilities of a port such as maximum frame size that can be received or end-to-end credit values. Again the N_Port login is performed by the initiator sending a PLOGI request frame and the target returning an accept.




Figure 23 PLOGI and Accept

Figure 23 shows the PLOGI and the Accept back from the target with their respective command codes. The payload of both the request and the accept PLOGI frames contain the service parameters of the initiator and the target.

Process Login (PRLI) is the process to allow two ports exchange service parameters relating to the FC-4 type they are using. Specifically what you would see here is SCSI type information being exchanged by both ports in a request and accept fashion as is shown in Figure 24.




Figure 24 PRLI and Accept

Some of the parameters exchanged during process login are whether the port is an initiator or target, transfer ready being used for reads or writes, and the FC4 type being used (0x08 for SCSI in these implementations).

Arbitration process

This section discusses the following topics:

- ◆ [“Overview” on page 43](#)
- ◆ [“Arbitration process steps” on page 44](#)
- ◆ [“Access fairness” on page 48](#)
- ◆ [“Opening and Closing the loop” on page 49](#)

Overview

Arbitration is the process that allows a loop port to gain access to the loop and ensure that only one port at a time is sending information. If this was not the case then ports could send data at the same time and interfere with each other. If two ports arbitrate at the same time to get on to the loop then there has to be a mechanism to handle multiple simultaneous requests. This is done by allowing the port with the lowest value AL_PA to have priority over ports with higher value AL_PAs. This could cause a situation where higher priority loop ports monopolize the loop, but to handle this a fairness mechanism is incorporated into the arbitration protocol (refer to [“Access fairness” on page 48](#)).

Before going into the arbitration process the concept of *Fill Words* needs to be understood. In Fibre Channel, even if there are no frames being transmitted, idle words are continuously transferred around the loop. When frames are being sent around the loop a certain amount of idles are also transmitted between these frames. These idles are a form of fill word and must exist between frames in a Fibre Channel environment. During arbitration it is necessary for a port to remove a fill word it has received and transmit a different fill word in its place. This is known as *fill word substitution*. When a port needs to transmit a fill word it will use the value contained in the current fill word. For example, if a port wants to start arbitrating on a loop which is not being used it needs to change the current fill word from ARB(x) to IDLE and once this is done the port can transmit ARB(x) instead of the received IDLEs.

Once the loop initialization phase has completed the loop is filled with IDLEs as each port is in the monitoring state. Once a port needs access to the loop it has to arbitrate to get this access. The six steps

involved are discussed next, using an example of a loop containing four ports with one port arbitrating as shown in [Figure 25](#).

Arbitration process steps

This section details the six arbitration process steps.

Arbitration Step 1




Figure 25 Arbitration Step 1

Step 1 of the arbitration process, shown in [Figure 25](#), shows a loop which is filled with idles and all ports in the monitoring status. The current fill word on all the ports is IDLE and thus received IDLEs are substituted with the current fill word which is IDLE on all the ports.

Arbitration Step 2




Figure 26 Arbitration Step 2

Step 2 (Figure 26) shows Port 1 arbitrating for access to the loop by changing its current fill word to ARB(1) and substituting ARB(1) for every received IDLE. If Port 1 receives any ARBs, it needs to check the priority of these ARBs and then either discard if they are a lower priority than ARB(1) or substitute if they are a higher priority.

Arbitration Step 3




Figure 27 Arbitration Step 3

In **Step 3** (Figure 27), once the ARB(1) is received by port 2 it changes its current fill word to ARB(1) and the ARB(1) is transmitted to the next port on the loop. The current fill word on the remaining ports is

also changed in this fashion allowing the ARB(1) to propagate around the loop.

Arbitration Step 4




Figure 28 Arbitration Step 4

In **Step 4** (Figure 28), once the ARB(1) is received by Port 1 it has won arbitration and immediately changes its current fill word to ARB(F0). Port 1 now will discard any received ARB(x) and thus this prevents any other loop port from winning arbitration as its ARB(x) could not make it around the loop.

Arbitration Step 5

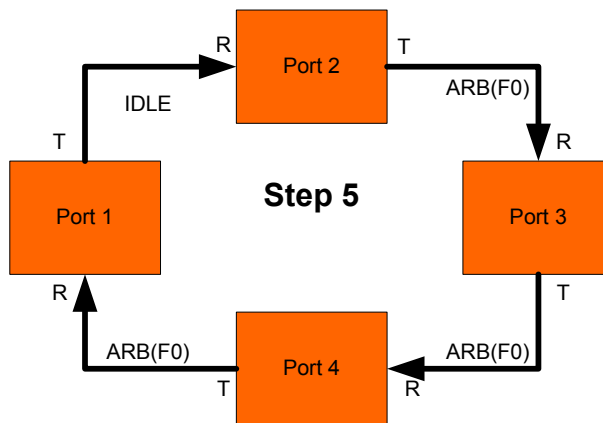


Figure 29 Arbitration Step 5

In **Step 5** (Figure 29 on page 46) each port that receives the ARB(F0) changes its current fill word to ARB(F0) and transmits the ARB(F0) whenever required. This is transmitted by the winner of the arbitration to determine if any other ports are arbitrating. If there are other ports arbitrating then they substitute their ARB(x) for the ARB(F0).

As long as no other port attempts to start arbitrating, the ARB(F0) makes it around the loop to port 1 and the current fill word on port 1 is changed to IDLE.

Arbitration Step 6




Figure 30 Arbitration Step 6

In **Step 6** (Figure 30) each port that receives the IDLE changes its current fill word to IDLE and as long as no other port is arbitrating the IDLE's travel around the loop. As long as Port 1 owns the loop it discards any received ARB(x) or IDLE and continues to send IDLE as its fill word. As it is discarding any received ARB(x) it prevents any other loop port from winning arbitration.

Figure 30 details how a single port arbitrating win ownership of the loop but one must also consider what happens if multiple ports begin arbitrating at the same time. This is handled by the fact that AL_PAs have different priorities and in this case the higher priority AL_PA (lower numerical value) will win the arbitration.

Access fairness

Arbitrated loop assigns different priorities to different AL_PA values and thus the situation could occur on a loop where the higher priority AL_PA ports could dominate the loop and prevent access from the lower AL_PA ports. To prevent this, **access fairness** is used to counteract AL_PA priority and ensure every loop port has an equal chance to access the loop. Access fairness does not mean how long a port that has won arbitration can use the loop, but rather ensures that once a loop port has won arbitration and given up control of the loop it will not arbitrate again until all other loop ports that are arbitrating have had an opportunity to access the loop. Ports that are not arbitrating do not affect the access fairness protocol. A loop port that follows the access fairness protocol is called a *fair* port and a port that does not is called an *unfair* port.

Access fairness applies in a certain time interval from when the first port on a loop wins arbitration to when no other port is arbitrating for ownership of the loop. This is controlled by each port on the loop setting a control variable called the *Access variable*. When a port wins arbitration it sets the access variable to 0 to ensure it cannot arbitrate again in the current fairness window until all other ports arbitrating have had a chance to win arbitration. The winning port knows what other ports are arbitrating as it is sending ARB(F0) and the ports that are trying to arbitrate are substituting their ARB(x) for this.

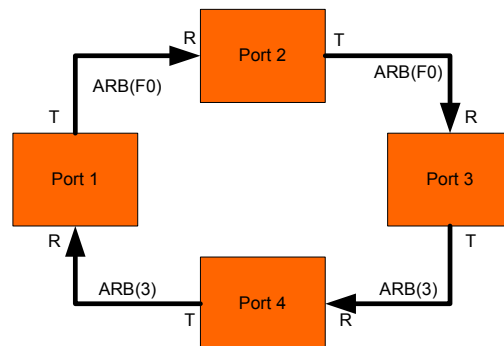


Figure 31 Access fairness window

If the winning port receives around an ARB(F0) it knows that no other port was arbitrating during the time that ARB(F0) was being sent around the loop which means the end of the current access fairness window. At the end of a fairness window all fair ports that

have their access variable set to 0 to prevent them arbitrating in that fairness window have this bit set again allowing them to arbitrate once more.

If the winning port does not receive back an ARB(F0), but instead an ARB(x) from a port trying to arbitrate, then it knows that another port is also arbitrating for access to the loop. Once the current winner relinquishes ownership of the loop it cannot arbitrate in the same fairness window so the other port arbitrating has a chance to win arbitration. This is shown in [Figure 31 on page 48](#) where Port 1 has won arbitration but Port 3 is also arbitrating for ownership of the loop. Access fairness will ensure that Port 3 will win arbitration once Port 1 has relinquished ownership.

Opening and Closing the loop

After a port has won arbitration it needs to select a destination port before it can send frames to it. It does this by opening a connection with the destination port by sending an open (OPN) ordered set with the AL_PA of the destination port. Once the destination port has received the OPN ordered set the loop circuit between both ports is open. At this stage all other ports on the loop behave as repeaters and allow frames to pass through on the way to the destination port.

Opening the loop

[Figure 32](#) shows Port 1 on a loop opening a circuit with Port 3 by sending an OPN ordered set with the destination port of 3 and once this is received Port 3 responds saying it was ready to receive a frame by sending back the R_RDY.

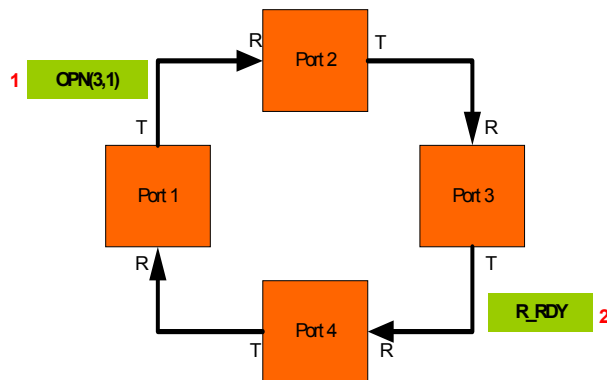


Figure 32 Opening a loop circuit

Figure 32 shows Port 1 on a loop opening a circuit with Port 3 by sending an OPN ordered set with the destination port of 3 and once this is received Port 3 responds saying it was ready to receive a frame by sending back the R_RDY.

As can be seen in Figure 33, there are different open ordered sets to open a circuit between ports.

Open Full-Duplex (Point-to-Point)	OPN(yx)	K28.5 D17.4 AL_PD AL_PS
Open Half-Duplex (Point-to-Point)	OPN(yy)	K28.5 D17.4 AL_PD AL_PD
Open Broadcast Replicate	OPN(fr)	K28.5 D17.4 D31.7 D31.7
Open Selective Replicate	OPN(yr)	K28.5 D21.4 AL_PD D31.7

Figure 33 Open primitive signals

- ◆ **OPN(yx)** is used to open a connection between two ports where the frame transmission will be in both directions. This operation allows both ports transmit at the same time and gives maximum utilization of the bandwidth since frame flow can be in both directions simultaneously.
- ◆ **OPN(yy)** is used when the originator of the open only wants to send frames to the destination port and wants to prevent the destination from sending frames to it other than link control frames such as Acknowledge, Busy, or Reject.
- ◆ **Open Replicate** is a mode ordered set used to implement multicast and broadcast capabilities in an arbitrated loop environment. The open ordered sets set the mode on the destination port to replicate which basically causes it to retransmit all received information and process the frames that have its address in the frame header.
 - **OPN(yr)** is used to selectively set replicate mode on a port specified by the AL_PD field of the ordered set.
 - **OPN(fr)** causes all ports on the loop to set replicate mode except the port transmitting the OPN(fr) ordered set.

Closing the loop

When all the frames are transmitted from one port to another then the ports may close the loop by sending a **Close** ordered set from one port to the other. When a port is finished transmitting frames it sends a CLS to the destination port and it enters the transmitted close state and waits for the other port to finish. When a port receives a CLS ordered set it enters the Received Close state finishes its frame transmission and sends a CLS. Once the port that initiated the close

receives back a CLS from the other port it either enters the monitoring state to allow another port win arbitration or it may send an open to establish a connection with another port.

The following three steps explain how to close a loop.

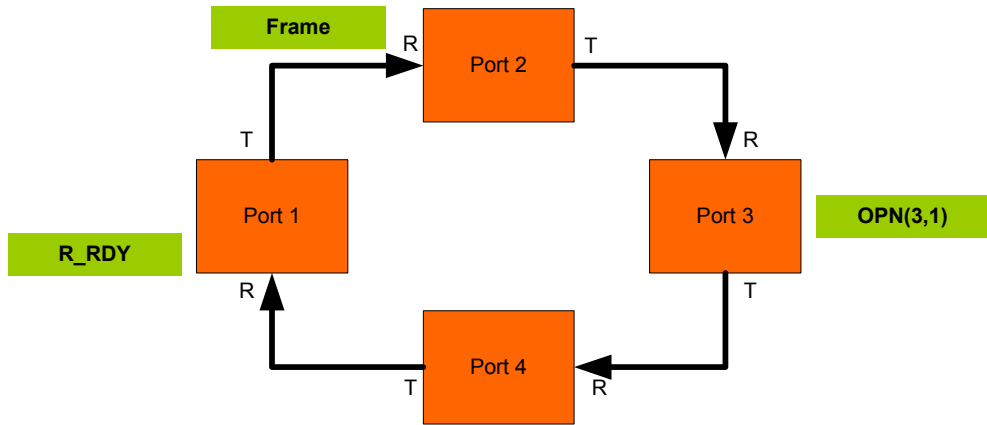



Figure 34 Close loop: Step 1

Step 1. Figure 34 shows an existing open loop between Port 1 and Port 3. Port 3 has received the OPN, responded by sending the R_RDY, Port 1 has received the R_RDY and is transmitting a frame to Port 3.



CFW = Current Fill Word

Figure 35 Close loop: Step 2

Step 2. Figure 35 on page 51 shows Port 1 sending a Close to Port 3 and Port 1 enters the transmitted close state. Port 3 can continue to transfer frames as long as it has credits to do so.

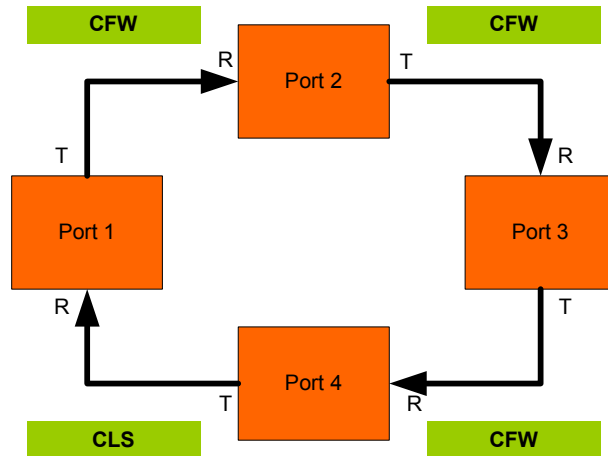


Figure 36 Close loop: Step 3

Step 3. When Port 3 receives the Close it enters the received close state, finishes its frame transmission and sends a CLS of its own (Figure 36). Port 3 also enters the monitoring state. When Port 1 receives the close it can enter the monitoring state or open a loop circuit with another port.

Alternate Buffer-to-Buffer Credit (BB_Credit)

In arbitrated loop the flow of frames between two ports is handled using a different credit model than the ones mentioned previously. Buffer-to-buffer credit and end-to-end credit flow control are both credit models that are negotiated between ports during the login process. For more information on buffer-to-buffer credit, refer to the "Buffer-to-Buffer credit (BB-Credit)" section in the *Networked Storage Concepts and Protocols TechBook*, available through the E-Lab Interoperability Navigator, **Topology Resource Center** tab, at <http://elabnavigator.EMC.com>.

For more information on end-to-end credit, the "End-to-End credit" section in the *Networked Storage Concepts and Protocols TechBook*, available through the E-Lab Interoperability Navigator, **Topology Resource Center** tab, at <http://elabnavigator.EMC.com>.

Alternate BB_Credit is used in arbitrated loop environments instead of login BB_Credit as each port on a loop may open up to 126 other ports on the loop and thus would have to keep track of the credit it has with each of the other ports on the loop from N_Port login time. This would be a significant complication. The Alternate BB_Credit model is an enhancement of the login BB_Credit model in that it begins with an initial login credit value which can be zero and can then be increased dynamically once a loop circuit between two ports is opened.

Before the ports log in, the available credit is assumed to be zero and even the login BB_Credit granted to another port may be zero. This login BB_Credit is used to set the available BB_Credit value when the

loop circuit is opened, although ports may or may not use this login BB_credit value.




Figure 37 Alternate credit

BB_Credit is then signalled dynamically by a port by sending one R_RDY for each available receive buffer. This allows ports that were opened with an available BB_Credit value of zero to begin frame transmission as soon as they receive an R_RDY from the receiving port. A loop port must indicate whether it supports login BB_Credit or Alternate Credit model. This is done during the N_Port login where the BB_Credit management bit (BBM) is either **on** for support of Alternate credit or **off** for support of the login BB_Credit model.

Figure 37 shows the situation where Port 1 and Port 3 during N_Port login indicated Alternate credit support and BB_Credit of zero. Thus when port 1 opens a loop circuit with port 3 by sending the OPN(3, 1) then port 1 has to wait until the R_RDY is received back from port 3 before a receive it can send a frame. The number of R_RDYs received by port 1 indicates the number of receive buffers available on port 3. Thus, when the loop is opened with an initial available BB_Credit value of zero there is a time delay in waiting for the first R_RDY to come back around where no data transfer can take place.

When the loop is closed the available BB_Credits are set to zero and any outstanding R_RDYs are not sent. The available BB_Credit value is then re-established when a loop circuit is opened between two ports.

This chapter contains information on QuickLoop.

- ◆ Overview 56
- ◆ Need for QuickLoop 57
- ◆ Defining QuickLoop 58
- ◆ Configuration examples 60

Overview

With data centers supporting HP-UX FC-AL and a variety of FC-SW types, EMC-supported solutions include the following:

- ◆ Single DS-16B operating in QuickLoop mode
- ◆ Two DS-16Bs participating in a QuickLoop *Partnership* (described under [“Concepts, terms, and rules” on page 58](#))
- ◆ Single DS-16B operating in mixed mode with QuickLoop and FC-SW nodes

Note: RPQs may be accepted for mixed-mode switches participating in a multi-switch fabric.

Need for QuickLoop

QuickLoop offers customers with HP-UX FC-AL hosts a migration path to full fabric environments. Using a Fibre Channel switch, FC-AL nodes can communicate with other private nodes over the fabric infrastructure.

EMC also supports QuickLoop with OpenVME hosts from ICL.

Defining QuickLoop

In a normal FC-AL hub environment, all traffic (for example: arbitrations, data requests, reads, and writes) is passed around the loop from node to node to node. The limitation is that only one node may send out traffic at a time. This solution does not scale well; all traffic on the loop is constrained to 100 MB/s total bandwidth.

QuickLoop allows hosts and storage to connect to a DS-16B using FC-AL drivers and HBAs.

Each port enabled for QuickLoop mode on the DS-16B is considered a *loople*t. If a host is on Port 0, one node is considered to be in that looplet. If an eight-port hub is attached to Port 0, there are seven nodes on that looplet. Data transfers between looplets are supported at 100 MB/s. With QuickLoop, there may be multiple data transfers capable of 100 MB/s each.

Concepts, terms, and rules

Note the following:

- ◆ A *Private Loop* device does not attempt to log in with a fabric and communicates only with other devices on the same loop.
- ◆ A *Public Loop* device (for example: HP-UX L-Class with A5158A HBA) logs in with a fabric and may communicate with both public and switched nodes.

Note: This is not a QuickLoop function, nor is it supported by EMC.

- ◆ QuickLoop nodes do not interact with fabric nodes.
FC-AL public nodes added to a QuickLoop effectively become private nodes. As a result, they cannot initiate communication with fabric nodes.
- ◆ EMC supports one QuickLoop per switch.
- ◆ QuickLoop may span a maximum of two switches. (This is known as a QuickLoop Partnership.) However, both switches may be part of only one QuickLoop.

A QuickLoop Partnership is supported in a standalone configuration. RPQs may be accepted for multiswitch fabrics with two of the switches participating in a QuickLoop Partnership.

- ◆ EMC supports the zoning of QuickLoop-enabled ports zoned by port number.
 - Since QuickLoop ports are hardware enforced, port zoning must be in effect.
 - You may have up to 16 of these zones in a single QuickLoop.
 - Port-zoning of QuickLoop enables ports isolates LIPs (loop initialization primitives).
 - For zoning requirements, refer to the following documents, available on [Powerlink](#):
 - *Topology Guide for DS-16B2 in a Symmetrix Environment*
 - *Departmental Switch Models DS-16B and DS-8B Zoning Reference Manual*
 - *DS-16B and DS-8B Topology Guide For Departmental Switches in a Symmetrix Environment*
- ◆ EMC supports a maximum QuickLoop of two DS-16Bs and two FC-AL hubs.
 - EMC supports a maximum of 16 HBAs in a single QuickLoop.
 - No hosts may be attached to a hub.
 - There is a maximum of eight Symmetrix Fibre Channel director ports on a hub.
- ◆ Supported firmware revisions for QuickLoop include 2.1.4a, 2.2.1a, 2.3.0, and 2.5.0d.
- ◆ Supported firmware revisions for port zoning of QuickLoop-enabled ports are 2.2.1a, 2.3.0, and 2.5.0d.
- ◆ Supported firmware revisions for single-switch mixed mode are 2.3.0 and 2.5.0d.
- ◆ When connecting to a QuickLoop enabled port, the Symmetrix Fibre Channel director port must be set to FC-AL.
- ◆ When set to FC-AL mode, a Symmetrix Fibre Channel director port will support private loop only.

Configuration examples

This section contains the following information:

- ◆ “Single QuickLoop configuration” on page 60
- ◆ “Dual QuickLoop configuration” on page 61
- ◆ “Mixed-mode configuration: Fabric and QuickLoop” on page 61
- ◆ “References” on page 62

Single QuickLoop configuration

Example: Single DS-16B operating in QuickLoop; HP K-Class host connected to a Symmetrix 8430, as shown in [Figure 38](#):



Figure 38 Single QuickLoop configuration

In an environment with only HP-UX FC-AL hosts, the switch may be 100% dedicated to a QuickLoop. The switch may only be part of single QuickLoop. Zoning of QuickLoop-enabled ports is not required because the entire switch is QuickLoop-enabled. With multi-initiators, however, zoning is recommended to restrict LIPs to that particular zone.

Dual QuickLoop configuration

Example: Two DS-16Bs operating in a QuickLoop Partnership, as shown in [Figure 39](#).




Figure 39 Dual QuickLoop configuration

Hosts and Symmetrix Fibre Channel director ports may be located anywhere on either switch. In a QuickLoop Partnership, there are two switches with one QuickLoop. However, there may be multiple-port zones of QuickLoop-enabled ports within a single QuickLoop.

Mixed-mode configuration: Fabric and QuickLoop

A mixed-mode switch environment refers to a DS-xxB switch that has both fabric nodes and QuickLoop ports, as shown in [Figure 40](#).




Figure 40 Fabric/QuickLoop mixed-mode configuration

Note: EMC supports only mixed mode with a single DS-xxB.

EMC requires that zoning of QuickLoop enabled ports be implemented whenever FC-SW and QuickLoop nodes are running on the same DS-xxB. The zone prevents LIPs from any particular zone from propagating into the fabric or into other zones.

In mixed mode, the switch is set in fabric mode with specific ports in QuickLoop mode.

EMC supports fabric zoning by World Wide Port Name (WWPN). Zoning with QuickLoop is hardware-enforced and therefore requires port zoning. Both WWPN and port zoning may be mixed in the same configuration.

References

For more information on latest qualified drivers and configurations, refer to the following:

- ◆ [E-Lab Navigator](#)
- ◆ *EMC Departmental Switches Model DS-16B2 Topology Guide for DS-16B2 in a Symmetrix Environment*, available on [Powerlink](#).

The fabric OS user's manual contains additional information on DS-xxB Telnet commands.

This chapter contains the following information on bridges:

- ◆ Overview 64
- ◆ SCSI-to-Fibre Channel bridges..... 65
- ◆ Crosspoint 4200 SCSI-to-fabric configuration..... 69
- ◆ ADIC SAN Gateway SCSI-to-fabric configuration 72
- ◆ ADIC SAN Gateway loop-to-fabric configuration..... 75

Overview

A bridge is a collection of both hardware and software whose primary function is facilitate communication between two dissimilar communication protocols. The bridge is the physical point were the dissimilar networks meet and the translation of protocols takes place. Bridges may also control traffic and security, filtering where necessary to boost SAN performance and contain sensitive data to particular areas of the SAN.

SCSI-to-Fibre Channel bridges

SCSI-to-Fibre Channel bridges allow an administrator to place SCSI and Fibre Channel devices into the same SAN environment. These devices communicate by using the bridges to translate commands between the two dissimilar protocols. Commercial bridges are available that support the connection of both target and initiator devices to the bridge's SCSI or Fibre Channel ports.

SCSI-to-Fibre Channel bridges can also be used to extend the distance between SCSI targets and initiator devices. For example, Fast-Wide Differential SCSI devices must be no farther than 25 meters from the initiator. A SCSI-to-Fibre Channel bridge allows extending the distance between the target and initiator to a total of 525 meters.

This section includes the following information:

- ◆ [“Operating modes” on page 65](#)
- ◆ [“Supported SCSI-to-Fibre Channel fabric bridges” on page 66](#)
- ◆ [“SCSI-to-fabric configuration envelope” on page 68](#)

Operating modes

This section discusses the storage and host modes.

Storage mode

SCSI-attached tape storage devices, tape library robotics and disk storage have been, and continue to be, a viable solution in the data storage environment. Whether these devices were legacy hardware, new purchases or advanced SCSI technology, they will also have to communicate with the Fibre Channel portion of the SAN environment.

In this mode of operation, such SCSI target devices as tape drives, tape libraries and disk drives are connected to the bridge's SCSI bus connections, and Fibre Channel initiators on the Fibre Channel link.

[Figure 41 on page 66](#) displays a simple storage mode configuration. This configuration, with the additional Fibre Channel switch support, can allow the incorporation of SCSI devices into the fabric.




Figure 41 Simple storage mode configuration

Host mode

You may find that some servers and operating systems will not support a Fibre Channel HBA. This topology is not qualified, and will be handled as required on a case-by-case basis.

SCSI-attached hosts, though supported by many bridge vendors, have not yet been validated in the EMC Quality Assurance process. The differences between Host and Storage mode should not be taken lightly and therefore should not be implemented without a thorough review of the customer's configuration and SAN's operational requirements.

Supported SCSI-to-Fibre Channel fabric bridges

EMC evaluates the reliability, functionality of performance of all products prior to their inclusion into our SAN solution sets. [Table 1](#) lists the products and their current feature support, as well as the EMC validation status. [E-Lab Navigator](#) lists the EMC-qualified features.

Table 1 Supported SCSI-to-Fibre Channel fabric bridges (page 1 of 2)

Feature	Crossroads CrossPoint 4200	ADIC SAN Gateway
Maximum Fiber Ports	1	6
Maximum SCSI Busses	4	4
Fibre Channel Single-Mode Ports	Yes	No
Fibre Channel Multimode Ports	Yes	Yes
Fibre Channel Switched Fabric	Yes	Yes

Table 1 Supported SCSI-to-Fibre Channel fabric bridges (page 2 of 2)

Feature	Crossroads CrossPoint 4200	ADIC SAN Gateway
Fibre Channel Arbitrated Loop	Yes	Yes
FibreAlliance MIB Compliance	Yes	Yes
SCSI Host Mode	Yes	Yes
SCSI Storage Mode	Yes	Yes
User-Initiated Diagnostics	Yes	Yes
Serial Communications Port	Yes	Yes
Ethernet Configuration Management	Yes	Yes
Web-Based Configuration Management	Yes	Yes
LAN-Free Backup	Yes	Yes
External LED Status Interface	Yes	Yes
Storage-Side LUN Masking	No	Yes
Storage-Side Fibre Channel/SCSI Channel Masking	No	Yes
Field-Upgradeable Firmware	Yes	Yes

About ADIC Technology, Inc.

ADIC Technology, Inc. manufactures the SAN Gateway, which can operate as a SCSI-to-Fibre Channel fabric bridge.

For more information on the ADIC SAN Gateway, SAN Router and value-added features unique to the ADIC SAN Gateway, visit www.adic.com.

About Crossroads Systems, Inc.

Crossroads Systems Inc. manufactures the Crosspoint 4200, which is validated to operate as SCSI-to FC-SW bridge. Crossroads also manufactures a newer series of SCSI-to-Fibre Channel bridges using the 4x50 family name.

It should be noted that many vendors may OEM the Crosspoint 4200 under a variety of different names (such as Brocade M Series EB1200, STK 3200, or Compaq FCTC II). The configuration guidelines below may be applied to these components as well as with the Crossroads 4x50 class of SAN appliances.

For more information on the Crossroads 4200, the 4x50 family of products and other value added features unique to the Crossroads 4200 and 4x50 family, visit www.crossroads.com.

SCSI-to-fabric configuration envelope

The following are guidelines and limitations for using the SCSI-to-Fibre Channel bridge:

- ◆ Only FC-SW configurations are supported for SCSI-to-Fibre Channel bridges.
- ◆ Only non-high-availability bridge configurations are supported.
- ◆ Bridges are currently supported only in a storage mode.
- ◆ Single-HBA zoning rules apply for HBA-to-bridge port zoning.
- ◆ ANSI-standard 100 [MB/s] Fibre Channel.
- ◆ Fibre Channel Distance up to 500 meters between Bridge and Switch port.
- ◆ Index (or Auto) Addressing scheme for SCSI-to-Fibre Channel device mapping.
- ◆ [E-Lab Navigator](#) lists the supported tape libraries and tape drives.
- ◆ For performance, a maximum of two high-bandwidth SCSI devices per bridge SCSI bus is recommended.

Crosspoint 4200 SCSI-to-fabric configuration

It should be noted that many vendors may OEM the Crosspoint 4200 under a variety of different names (such as Brocade M Series EB1200, STK 3200, or Compaq FCTC II). The configuration guidelines below may be applied to these components as well as with the Crossroads 4x50 class of SAN appliances.

To configure a SCSI tape drive or tape library device into the SAN environment using the Crossroads Crosspoint 4200 as a SCSI-to-Fabric Bridge, follow the steps outlined in the following paragraphs.

Initial setup

Follow the Crossroads procedures for the initial setup and power-on of the bridge. If you intend to configure the Crosspoint 4200 over the Web interface, make sure that you have properly configured the unit's IP address, netmask and TCP/IP gateway address. This must be initially done through the serial interface, using a PC and a terminal server program. You may also use the PC and Telnet into the bridge to complete the configuration.

1. Set SCSI port characteristics — Prior to the connection of a SCSI tape devices (targets), use the Crossroads procedures to set the mode of operation. Using Fibre Channel hosts to communicate with SCSI devices is referred to as Initiator Mode. Initiator Mode is the default operation mode for the Crosspoint 4200.
2. Set fibre port characteristics — The Crosspoint 4200's fiber connection can automatically detect if it is connected to a fabric and configure the port for fabric operation without intervention. If this is not the initial configuration for the bridge, Telnet into the bridge to verify the status of the configuration before proceeding.
3. Employ value-added features — At this point you may also decide to start to employ some of the value added features supplied by Crossroads. One of interest would be Fibre Channel-to-SCSI mapping. This allows you to manually configure how the physical target device LUNs will be seen on the HBAs connected to the bridge.

Refer to your current SAN/ESN management suite prior to enabling and configuring any additional management options available on the Crosspoint 4200. Any redundancy in this area may result in unexplained outages and difficulty in troubleshooting SAN issues in the future.

Target connections

It is recommended when installing new SCSI devices into the Crosspoint 4200 that you start with both the bridge and the SCSI target devices (Tapes or Disks) powered off. The Crosspoint 4200 will assign new LUN numbers and present them to the Hosts based on how they are discovered. With this in mind, connect the SCSI devices into the Crosspoint 4200's SCSI busses as you would like them to be numbered.

After the SCSI connections have been made, you may power on SCSI target devices; then power on the Crosspoint 4200.

At this point you may also decide to start to employ some of the value-added features supplied by Crossroads. One of interest would be Fibre Channel to SCSI Mapping. This allows you to manually configure how the physical target device LUNs will be seen on the HBAs connected to the bridge.

Refer to your current SAN/ESN management suite prior to enabling and configuring any additional management options available on the Crosspoint 4200. Any redundancy in this area may result in unexplained outages and difficulty in troubleshooting SAN issues in the future.

If SCSI devices are moved, added, or removed, you may need to reboot the bridges for the changes to take effect.

Remember to record the Crossroads, assigned LUN numbers, because they will be needed later when you perform any specific HBA-to-LUN discovery, mapping and persistent binding.

Bridge-to-SAN connections

The fiber port on the Crosspoint 4200 is provided a unique World Wide Port Name (WWPN) by the Crosspoint 4200. You should verify the current WWPN currently configured on the bridge and record this future use during the fabric zoning stage.

Host connections and configuration

As in any other SAN configuration, you will connect the host HBAs directly into the fabric backbone. Once both the host and Crossroads, Crosspoint 4200 are visible to the fabric you may start the procedures for zoning and the distribution of devices across the SAN.

At this point you may also want to perform any host specific procedures necessary to discover or map the new devices LUNs that will be discovered at the completion of the configuration.

Note: Refer to the appropriate host-specific configuration information in the EMC Host Connectivity Guides.

Fabric zoning

For the Crossroads, Crosspoint 4200 to communicate with the host's HBA, follow the current supported practices for HBA and bridge WWPN zoning on your fabric switch.

Final setup

Once you have completed the fabric zoning, you should return to the Crosspoint 4200 to ensure that the server's HBA can communicate with the Crosspoint 4200. After communication has been validated you should reboot your host machines so that it can discover the list of new SCSI devices.

ADIC SAN Gateway SCSI-to-fabric configuration

To configure a SCSI tape drive or tape library device into the SAN environment using the ADIC SAN Gateway as a SCSI-to-fabric bridge, follow the steps outlined in this section.

Initial setup

Follow ADIC's documented procedures for the initial setup and power-on of the Gateway. If you intend to configure the SAN Gateway through the Web interface, make sure that you have properly configured the unit's IP address, Netmask and TCP/IP Gateway address. You must also install the client/server portion of the Web software somewhere on your IP network.

Initial configuration of the IP interface must be done through the serial interface, using a PC and a terminal server program. You may also use a PC and telnet into the SAN Gateway to complete the configuration.

1. Set SCSI port characteristics — Prior to the connection of a SCSI tape devices (targets), use the ADIC procedures to set the Channel Mode of the SCSI port on the SAN Gateway in use to **Initiator**. The correct port designation can easily be remembered by always configuring the port to act as the opposite of the devices that you are installing. Since you are adding a target you need to configure the SAN Gateway port to act as an initiator.
2. Set fibre port characteristics — Before you connect the SAN Gateway into the fabric you will need to set the characteristics of the fiber port. Use the ADIC procedures for setting the Port Mode of the fiber connection to **Target**. This can be remembered in the same fashion as discussed above by always configuring the port to act as the opposite of the devices that you are installing.

Since you are going to be communicating with a host initiator device the SAN Gateway port must be seen as a target device. Lastly, since you are connecting to a fabric switch, you will want to set the Connection Options to **Point-to-Point** mode.

3. Employ value-added features — At this point you may also decide to start to employ some of the value-added features supplied by ADIC. EMC recommends that you review and possibly employ Channel Zoning and Virtual Private SAN.

The Channel Zoning option is currently a free offering that allows you to designate which SCSI channels can communicate with which Fibre Channels. You may choose to allow total communication between the different channels or select any combination of individual target-to-initiator paths.

The Virtual Private SAN (VPS) option requires a license and is used to designate which device LUNS can be seen on which individual HBAs. Once enabled you will have to either install a supplied software package on your host machines to register the HBA information with the SAN Gateway or manually configure the information into the Gateway. Refer to the ADIC User Manual for additional features and procedures associated with the Virtual Private SAN option.

It should be noted that these features may already be supplied in your current SAN/ESN management tools. Refer to your current SAN/ESN management suite prior to enabling and configuring any additional management options. Any redundancy in this area may result in unexplained outages and difficulty in troubleshooting SAN issues in the future.

Target connections

It is recommended when installing new SCSI devices into the SAN Gateway that you start with both the Gateway and the SCSI target devices powered off. The SAN Gateway will assign new LUN numbers and present them to the hosts based on how they are discovered. With this in mind, connect the SCSI devices into the SAN Gateway's SCSI busses as you would like them to be numbered.

After the SCSI connections have been made, power on the SCSI target devices. You may then power on the SAN Gateway.

If a SCSI device is moved, added, or removed, you can run the SCSI **rescan** command on the Gateway to get the correct device picture.

Remember to record the ADIC-assigned LUN numbers, because they will be needed later when you perform any specific HBA-to-LUN discovery, mapping and persistent binding.

Bridge-to-SAN connections

Each fiber port on the SAN Gateway is provided a unique World Wide Port Name (WWPN) by the SAN Gateway. The Gateway does this by:

1. Changing the first pair of digits in the WWN from **10** to **20**.
2. Substituting the next two digits in the WWN into a two-digit representation of the fiber port number being used.

For example, if you had connected fiber port 2 on the SAN Gateway (WWN 10:00:00:60:45:16:0A:0D) into the fabric switch, the WWPN that would appear on the fabric switch would be 20:02:00:60:45:16:0A:0D.

Host connections and configuration

As in any other SAN configuration you will connect the host HBAs directly into the fabric backbone. Once both the host and ADIC SAN Gateway are visible to the fabric you can start the procedures for zoning and the distribution of devices across the SAN.

At this point you may also want to perform any host-specific procedures necessary to discover or map the new devices LUNs that will be discovered at the completion of the configuration.

Note: Refer to the appropriate host-specific configuration information in the EMC Host Connectivity Guides.

Fabric zoning

For the ADIC SAN Gateway to communicate with the host's HBA, follow the current supported practices for HBA and Gateway WWPN zoning on your fabric switch.

Final setup

Once you have completed the fabric zoning, you should return to the SAN Gateway to ensure that the server's HBA can communicate with the SAN Gateway. After communication has been validated you should reboot your host machines so that it can discover the list of new SCSI devices.

ADIC SAN Gateway loop-to-fabric configuration

To configure an FC-AL tape drive into the SAN environment using the ADIC SAN Gateways as a loop-to-fabric bridge, follow the steps outlined in this section.

Initial setup

Follow ADIC's documented procedures for the initial setup and power-on of the Gateway. If you intend to configure the SAN Gateway through the Web interface, make sure that you have properly configured the unit's IP address, Netmask, and TCP/IP Gateway address. You must also install the client/server portion of the Web software somewhere on your IP network.

Initial configuration of the IP interface must be done through the serial interface, using a PC and a terminal server program. You may also use a PC and Telnet into the SAN Gateway to complete the configuration.

1. Set fibre port characteristics as a target — Prior to the connection of an FC-AL tape device (Target), use the ADIC procedures to set the **Fibre Channel Mode** of the Fibre Channel port on the SAN Gateway in use to **Public Initiator** and the **Connection Options** to **Loop only**. The correct port designation can easily be remembered by always configuring the port to act as the opposite of the devices you are installing. Since you are adding a target you need to configure the SAN Gateway port to act as an initiator.
2. Set fibre port characteristics as link connection — Before you connect the SAN Gateway into the fabric you will need to set the characteristics of the fiber port. Use the ADIC procedures for setting the **Port Mode** of the fiber connection to **Target**. This can be remembered in the same fashion as discussed above by always configuring the port to act as the opposite of the devices that you are installing.

Since you are going to be communicating with a host initiator device the SAN Gateway port must be seen as a target device. Lastly, since you are connecting to a fabric switch, you will want to set the **Connection Options** to **Point to Point** mode.

3. Employ value-added features — At this point you may also decide to start to employ some of the value-added features supplied by ADIC. EMC recommends that you review and possibly employ Channel Zoning and Virtual Private SAN.

The Channel Zoning option is currently a free offering that allows you to designate which SCSI channels and Fibre Channels can communicate with which Fibre Channels. You may choose to allow total communication between the different channels or select any combination of individual target-to-initiator paths.

The Virtual Private SAN (VPS) option requires a license and is used to designate which device LUNS can be seen on which individual HBAs. Once enabled you will have to either install a supplied software package on your host machines to register the HBA information with the SAN Gateway or manually configure the information into the Gateway. Refer to the ADIC User Manual for additional features and procedures associated with the Virtual Private SAN option.

Note that these features may already be supplied in your current SAN/ESN management tools. Refer to your SAN/ESN management suite prior to enabling and configuring any additional management options. Any redundancy in this area may result in unexplained outages and difficulty in troubleshooting SAN issues in the future.

Target connections

It is recommended when installing a new FC-AL devices into the SAN Gateway that you start with both the Gateway and the FC-AL target devices powered off. The SAN Gateway will assign new LUN numbers and present them to the hosts based on how they are discovered. With this in mind, connect the FC-AL device into the SAN Gateway's Fibre Channel port as you would like it to be numbered.

After the loop connections have been made, power on the FC-AL target devices. You may then power on the SAN Gateway.

Remember to record the ADIC-assigned LUN numbers, because they will be needed later when you perform any specific HBA-to-LUN discovery, mapping and persistent binding.

Bridge-to-SAN connections

Each fiber port on the SAN Gateway is provided a unique World Wide Port Name (WWPN) by the SAN Gateway. The Gateway does this by:

1. Changing the first pair of digits in the WWN from **10** to **20**.
2. Substituting the next two digits in the WWN into a two-digit representation of the fiber port number being used.

For example, if you had connected fiber port 2 on the SAN Gateway (WWN 10:00:00:60:45:16:0A:0D) into the fabric switch, the WWPN that would appear on the fabric switch would be 20:02:00:60:45:16:0A:0D.

Host connections and configuration

As in any other SAN configuration you will connect the host HBAs directly into the fabric backbone. Once both the host and ADIC SAN Gateway are visible to the fabric you can start the procedures for zoning and the distribution of devices across the SAN.

At this point you may also want to perform any host specific procedures necessary to discover or map the new devices LUNs that will be discovered at the completion of the configuration.

Note: Refer to the appropriate host-specific configuration information in the EMC Host Connectivity Guides.

Fabric zoning

For the ADIC SAN Gateway to communicate with the host's HBA, follow the current supported practices for HBA and Gateway WWPN zoning on your fabric switch.

Final setup

Once you have completed the fabric zoning, you should return to the SAN Gateway to ensure that the server's HBA can communicate with the SAN Gateway. After communication has been validated you should reboot your host machines so that it can discover the list of new FC-AL devices.

Reference

[Chapter 1, "Fibre Channel Arbitrated Loop \(FC-AL\),"](#) contains more information on connection FC-AL tapes into a switched fabric.

Interfacing Arbitrated Loop to Switched Fabric

This chapter contains the following information for interfacing arbitrated loop to switched fabric:

- ◆ Overview 80
- ◆ Operating modes..... 81
- ◆ Connectivity devices that support FC-AL..... 82

Overview

Loop-to-fabric Fibre Channel bridges allow an administrator to place FC-AL and FC-SW devices into the same SAN environment. These devices communicate by using the bridges to translate commands between the dissimilar protocols. Fibre Channel switches and bridges are available that support FC-AL to FC-SW protocol translation. [E-Lab Navigator](#) lists the supported connectivity devices that allow the connection of FC-AL initiators.

Operating modes

This section briefly discusses the storage and host modes.

Storage mode

In *storage mode* operation, such target devices as tape drives and disk drives are connected to the bridge's Fibre Channel ports. Other Fibre Channel ports on the bridge are used to connect the bridge into the fabric. The number of fabric connections from bridge to fabric depends on the actual bandwidth needed to consistently stream data to the target devices on the bridge.

Host mode

In *host mode* operation, initiator devices are attached to either a SCSI port or Fibre Channel port on the bridging device. The bridge is then connected through another Fibre Channel port to the fabric.

You may find that some servers and operating systems will not support Fibre Channel in the chassis or FC-SW settings on an HBA. This topology is not qualified, and will be handled as required on a case-by-case basis.

Connectivity devices that support FC-AL

The EMC Fibre Channel implementation supports these Fibre Channel connectivity devices for tape:

- ◆ EMC Products:
 - Connectrix DS-16B (Brocade SilkWorm 2xxx)
 - Connectrix DS-8B (Brocade SilkWorm 2xxx)
 - Connectrix DS-16B2 (Brocade SilkWorm 3800)
 - Connectrix DS-24M2 (Brocade M Series 4500)
- ◆ Brocade Products:
 - SilkWorm 2xxx series (EMC Connectrix DS-16B/8B)
 - SilkWorm 3800/3200 (EMC Connectrix DS-16B2/DS-8B2)
- ◆ Brocade M Series Products
 - Brocade M Series 4500 (EMC Connectrix DS-24M2)
- ◆ ADIC Products: SAN Gateway

This section contains information on the following:

- ◆ ["Connectrix DS-16B, DS-16B2 \(Brocade SilkWorm Series\)" on page 82](#)
- ◆ ["Brocade M Series ES-1000" on page 85](#)
- ◆ ["ADIC SAN Gateway" on page 87](#)
- ◆ ["Interfacing arbitrated loop to switched fabric summary" on page 89](#)
- ◆ ["Loop-to-fabric configuration envelope" on page 90](#)

Connectrix DS-16B, DS-16B2 (Brocade SilkWorm Series)

These switches all provide the same functionality. Each port on these switches is capable of auto-negotiating and auto-initializing the port for public loop or fabric operation. Auto-configuration occurs immediately after the devices are connected to the switch port. At the completion of the auto-configuration phase (almost instantaneously), the port should appear as an FL_Port in the name server list on the switch. No additional software or licensing is required for this functionality.

Note: Do not configure these ports as QuickLoop/Fabric Assist ports. The QuickLoop/Fabric Assist mechanisms are not required for public loop device support.

When the negotiation and initialization is complete, you will also be able to view the device's WWPN in the switch's name server list.

Port count Each DS-16B and DS-16B2 (Brocade SilkWorm) has 16, and each DS-8B has 8, Fibre Channel ports that individually and automatically negotiate whether to communicate using Fibre Channel arbitrated loop or switched fabric protocols.

Features The Connectrix DS-16B/8B and DS-16B2 are Departmental Switches that can provide all the Fibre Channel protocol features associated with a fabric switch. ISLs may be connected to either another Brocade switch (in native Brocade Fabric mode) or to a Brocade M Series switch (using Open Fabric mode). Using Connectrix DS-16B/B2 in Open Fabric mode allows you to easily connect FC-AL devices into a Brocade M Series fabric.

The ability of the Connectrix DS-16B/8B/16B2 to automatically negotiate the correct protocol on an individual port basis allows you to mix E_Ports, N_Ports, and NL_Ports on the same switch.

Fabric Since the Connectrix DS-16B/8B/16B2 is a switch, it requires a unique Domain ID before connecting into an existing fabric. EMC provides recommendations for fabric topology configurations and sizes that must be followed when using a switch to incorporate FC-AL tapes into a SAN environment.

The Connectrix DS-16B/8B/16B2 can be used in either a homogeneous Brocade fabric or a heterogeneous fabric that includes Brocade M Series switches. When incorporating a Brocade switch into a fabric that includes Brocade M Series switches, the fabric mode on all switches must be set to Open Fabric.

Maintenance Each tape drive is individually attached to the DS-16B/8B/16B2, and forms a single-device loop with the switch port. This allows the user to work individually on any tape device without experiencing an interruption associated with a Loop Initialization Protocol (LIP) exchange across the other devices.

Note: The tape library robot may be attached over either Fibre Channel or SCSI. Be sure to understand your specific connectivity requirements for your library. ("[SCSI-to-Fibre Channel bridges](#)" on page 65 provides more information.)

Zoning Public loop and switched tapes attached to a switch register their unique WWNs with the switch's name server. This allows the user to directly zone the tape's WWN with the WWN of the initiator device (HBA).

For more information, refer to the "Sharing tape and disk on the same HBA" section in the *Backup and Recovery in a SAN TechBook*, available through the E-Lab Interoperability Navigator, **Topology Resource Center** tab, at <http://elabnavigator.EMC.com>.

SAN configuration recommendations

EMC recommends that you first identify your actual data throughput requirements prior to locking down any configuration. As a best practice, start by taking the theoretical throughput of the switch's Fibre Channel connectivity capabilities (for example: 1 Gb/s, 2 Gb/s, and so on) and divide this by the tape drive's theoretical throughput capability (not its Fibre Channel connector specifications) to identify how many tape drives can be streaming simultaneously on a single ISL. For example:

- ◆ Assume a two switch fabric (A and B).
- ◆ Each tape drive in the library can support a throughput of 25 MB/s.
- ◆ Both switches in the fabric have 100 MB (half-duplex) switch ports.
- ◆ $\text{MB per ISL} / \text{MB per tape drive} = \text{tapes per ISL (rounded down)}$.
- ◆ $100 \text{ MB} / 25 \text{ MB} = 4 \text{ tape drives}$.

For particular environments that do not require the same simultaneous tape access, you may be able to reduce the required ISLs and increase the number of tape drives per switch. You may also be able to use the backup application's drive pooling facility to distribute your simultaneous drive load across multiple ISLs.

[Figure 42 on page 85](#) shows two tape pools, each on its own core switch. Placing the pools on separate switches guarantees that different ISLs are used when each pool is accessed. You can also increase the overall availability of tape resources by spreading them across the fabric.




Figure 42 Examples of tape pools

Brocade M Series ES-1000

Port count Each Brocade M Series ES-1000 has eight loop-only Fibre Channel ports acting as an 8-port hub, as well as a single Fibre Channel B_Port (Bridge Port) used for attaching the ES-1000 to other Brocade M Series Fibre Channel switches.

- Features**
- ◆ The ES-1000 product is fully manageable by EMC Connectrix Manager.
 - ◆ Each device attached to the ES-1000 registers its WWN with the switch's name server, allowing the individual devices to be zoned independently.
 - ◆ All devices attached to the 8-port hub will share the bandwidth of the 1 Gb B_Port ISL to the remainder of the fabric.

- ◆ Not all industry switches support the B_Port protocol; for this reason the ES-1000 may be connected only to ED-64M, DS-16M, DS-32M, ED-1032.
- ◆ Current versions of the ES-1000 firmware offer support for Open Fabric mode operation in a heterogeneous environment. It should be noted that the ES-1000 must still be directly connected to the Brocade M Series family of switches.

Fabric

Since the ES-1000 is a workgroup switch and supports B_Ports (specialized, limited functionality E_Ports), it requires a unique Domain ID before connecting into an existing fabric. EMC provides recommendations for fabric topology configurations and sizes that must be followed when using a switch to incorporate tapes into a SAN environment.

It should be noted that since the ES-1000 provides only one B_Port, the loss of this connection would constitute a single point of failure for the fabric.

Users may use the ES-1000 in either a homogeneous Brocade M Series fabric or a heterogeneous fabric, as long as the ES-1000 is attached to a Brocade M Series switch. Users should also familiarize themselves with the requirements and operation of an heterogeneous fabric before making the decision to include a Brocade switch in a Brocade M Series switch fabric.

Maintenance

The ES-1000 forms a Fibre Channel loop consisting of the entire set of devices that are connected to the FC-AL ports. Since all devices are in the same loop, maintenance on any one device may cause a LIP (loop initialization protocol) message to be set to all devices on the loop. LIPs may cause interruptions in backup jobs.

Zoning

Since devices attached to the switch register their own unique World Wide Names (WWN) with the switch's name server, the user will be able to directly zone the individual tape drives or other target devices with the WWN of the initiator device (HBA).

SAN configuration recommendations

EMC recommends that you first identify your actual data throughput requirements prior to locking down any configuration. As a best practice, start by taking the theoretical throughput of the of the switch's Fibre Channel connectivity capabilities (for example, 1 G) and divide this by the tape drive's theoretical throughput capability (not its Fibre Channel connector specifications) to identify how many tape drives can be streaming simultaneously on a single ISL.

Since the Brocade M Series ES-1000 has only a single ISL, the simultaneous access requirements to the tape drives may need to be relaxed when fully populating the switch with tape drives. Users may also decide to distribute their tape devices across multiple ES-1000 devices to limit the impact of an ISL or ES-1000 link failure.

ADIC SAN Gateway

Port count Each ADIC SAN Gateway has up to six Fibre Channel ports. Each Fibre Channel port can be manually configured for either FC-AL or FC-SW, or can be configured to detect and configure itself based on the current attachment.

The SAN Gateway may also be configured with up to four additional SCSI adapters.

Features The ADIC SAN Gateway comes with its own client/server management application that may be used to configure, assign host-to-device access, or assign devices to specific channels. A single-client/server setup may be used to manage multiple gateways simultaneously.

Fabric Since the ADIC SAN Gateway is not a switch, it may be added to any size fabric without the need for a unique fabric Domain ID. Depending on the individual Fibre Channel port setting, the ADIC SAN Gateway connects to the switch and logs in to the switch using the FC-SW protocols. When the ADIC SAN Gateway logs in to the switch it registers the WWN of the port of the SAN Gateway that is being used. All Fibre Channel ports on the SAN Gateway have the same lower WWN. Only the second pair of the upper WWN is changed to indicate the local Fibre Channel port of the gateway.

EMC provides recommendations for fabric topology configurations and sizes that must be followed when using a SAN Data Gateway to incorporate tapes into a SAN environment.

Maintenance Each FC-AL tape drive is individually attached to a Fibre Channel port on the SAN Gateway, and forms a single-device loop with the gateway. This allows the user to work individually on any tape device without experiencing an interruption associated with a Loop Initialization Protocol exchange across the other devices.

Zoning Devices attached to the gateway are automatically assigned a unique LUN by the gateway. The user may use the management application

to change this number, but EMC recommends that the configuration list a tape library first with the lowest LUN, followed by its tape drives listed in sequential order. This ordering can be repeated for multiple tape libraries on the same gateway.

A feature on the gateway called *channel zoning* allows the user to associate target devices with specific gateway Fibre Channel ports. Once the gateway's Fibre Channel port(s) are connected to a switch in the fabric, the gateway registers its WWN(s) as stated above.

Zoning is performed using the fabric's management environment (for example, ESN Manager). A zone will consist of the server's HBA WWN and the gateway Fibre Channel port WWN. Devices will appear to the host in a similar fashion to the way disk devices are displayed behind a SCSI controller (that is, each tape device will be referenced under a target and LUN combination). Unique LUNs will be assigned for each tape driver or library robot. The target number may differ depending on how the devices were spread across the Fibre Channel ports on the SAN Gateway.)

SAN configuration recommendations

EMC recommends that you first identify your actual data throughput requirements prior to locking down any configuration. As a best practice, start by taking the theoretical throughput of the of the gateway's Fibre Channel connectivity capabilities (for example, 1 G) and divide this by the tape drive's theoretical throughput capability (not its Fibre Channel connector specifications) to identify how many tape drives can be streaming simultaneously on a single Fibre Channel connection from the gateway to the switch.

For particular environments that do not require the same simultaneous tape access, you may be able to reduce the required ISLs and increase the number of tape drives per gateway.

Note: Multiple links between the SAN Gateway and the fabric are possible. To further increase the availability you may also attach links to separate fabrics; however, this lowers the number of tape drives that can be attached.

Configuration guidelines

"[Loop-to-fabric configuration envelope](#)" on page 90 provides some configuration guidelines.

Interfacing arbitrated loop to switched fabric summary

[E-Lab Navigator](#) lists the EMC-qualified features.

Table 2 Interfacing arbitrated loop to switched fabric: Summary (page 1 of 2)

Feature	Connectrix DS-16B/8B, DS-16B2	Brocade SilkWorm	ADIC SAN Gateway	Brocade M Series ES-1000
Max Fiber Ports	16 FC-AL/FC-SW (DS-8B has 8)	8 or 16 FC-AL/FC-SW	6 FC-AL/FC-SW	8 FC-AL, 1 FC-SW
Max SCSI Busses	None	None	4	None
Fibre Channel Single-Mode Ports	Yes	Yes	Yes	Yes
Fibre Channel Multimode Ports	Yes	Yes	Yes	Yes
Fibre Channel Switched Fabric	Yes	Yes	Yes	Yes
Fibre Channel Arbitrated Loop	Yes	Yes	Yes	Yes
FibreAlliance MIB Compliance	Yes	Yes	Yes	Yes
SCSI Host Mode	No	No	Yes	No
SCSI Storage Mode	No	No	Yes	No
User-Initiated Diagnostics	Yes	Yes	Yes	Yes
Serial Communications Port	Yes, DS-8B only	Yes, 8-port switch only	Yes	Yes
Configuration Management	<ul style="list-style-type: none"> • Web Tools • ESN Manager • Telnet • Fabric Manager 	<ul style="list-style-type: none"> • Web Tools • ESN Manager • Telnet • Fabric Manager 	<ul style="list-style-type: none"> • SAN Director • HyperTerminal 	<ul style="list-style-type: none"> • EFCM • ESNM • Telnet
LAN-Free Backup	No; however, you can back up the switch/zoning configuration	No; however, you can back up the switch/zoning configuration	Yes	No
External LED Status Interface	Yes	Yes	Yes	Yes

Table 2 Interfacing arbitrated loop to switched fabric: Summary (page 2 of 2)

Feature	Connectrix DS-16B/8B, DS-16B2	Brocade SilkWorm	ADIC SAN Gateway	Brocade M Series ES-1000
Storage-Side LUN Masking	No	No	Yes	No
Storage-Side Fibre Channel/SCSI Channel Masking	N/A	N/A	Yes	N/A
Field-Upgradeable Firmware	Yes (disruptive)	Yes (disruptive)	Yes	Yes

Loop-to-fabric configuration envelope

The following are guidelines and limitations for using the loop-to-fabric bridge:

- ◆ Only non-high-availability bridge configurations are supported.
- ◆ Single-HBA zoning rules apply for HBA-to-bridge port zoning.
- ◆ ANSI-standard 100 (MB/s) Fibre Channel.
- ◆ Fibre Channel Distance up to 500 meters between bridge and switch port.

[E-Lab Navigator](#) lists the supported tape drives.

Storage Area Network Management

This chapter contains information on SAN management.

- ◆ Distance topology 92
- ◆ Capacity topology in the loop environment 93
- ◆ Consolidation topology in the arbitrated loop environment 94
- ◆ Combined topologies 95

Distance topology

In the early years of ESCON, both multimode for short distances and single mode for long distances were offered. However, it was found that most data centers did not want their long distance links to terminate at their mainframes or storage equipment. For both availability and resource management reasons, they preferred to terminate the long distance links on an ESCON director, and then implement a connection to another ESCON director.




Figure 43 FC-AL high-availability distance topology example

By replicating the hubs (as shown in [Figure 43](#)), as well as the server and storage connections at each site, a highly fault-resistant topology ensures maximum availability for the site-to-site interconnection.

Capacity topology in the loop environment

By using a hub, the host system includes a fault-resilient connection to a large pool of Symmetrix storage, reducing server I/O slot count.




Figure 44 FC-AL capacity expansion topology example

HP-UX PV Links play an essential role in providing high availability in many Fibre topologies. EMC PowerPath® does the same when deployed with Fibre Channel. These host-based software products provide availability during the failure and repair of one of the hubs. In addition, PowerPath provides load balancing services, which facilitates performance management of the system.

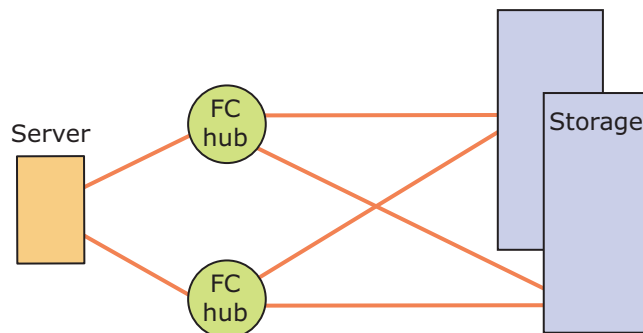


Figure 45 FC-AL high-availability capacity expansion topology example

Consolidation topology in the arbitrated loop environment

Figure 46 and Figure 47 are examples of the consolidation topology in a loop environment.




Figure 46 FC-AL consolidation topology example

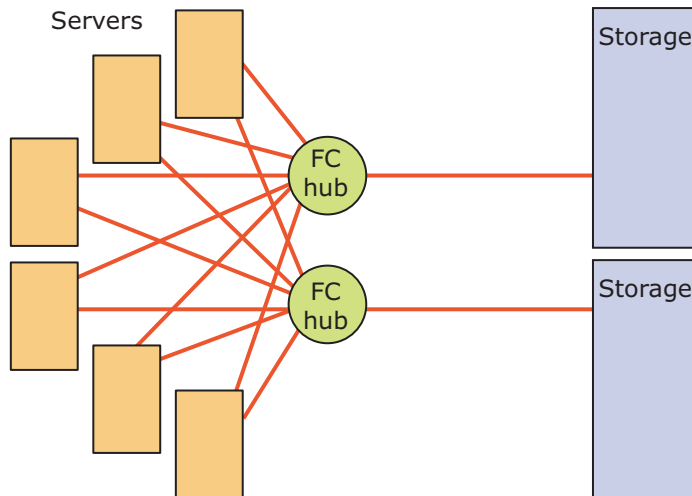


Figure 47 FC-AL high-availability consolidation topology example

Figure 47 shows six hosts, each with two FC-AL ports, which require a total of twelve storage connections. By using a pair of Fibre Channel hubs, all twelve server connections can be consolidated into two storage connections. Servers with low I/O demands (for example, they may be bottlenecked on CPU or networking resources) can share the storage bandwidth of two Symmetrix connections.

Combined topologies

Topologies can be combined for maximum efficiency, achieving large Symmetrix storage capacity for many servers while minimizing the necessary number of host bus adapters and Symmetrix Fibre Channel ports.

The three basic topologies can be combined to take advantage of the benefits of each.




Figure 48 FC-AL combined hub topologies example

This chapter contains the following information on CNT (Inrange):

- ◆ [Configuring CNT \(Inrange\).....](#) 98
- ◆ [IOCP considerations.....](#) 99

Configuring CNT (Inrange)

This section lists the support product and notes on topology support.

Supported product

FC-9000

Topology support

Note the following:

- ◆ Up to two domains and a single hop are supported.
- ◆ Intermixing FICON and Symmetrix-FCP is not supported.
- ◆ Intermixing FICON and SRDF is not supported.
- ◆ CUP is supported.
- ◆ The switch address range is **x'01'** to **x'EF** (1-239).

Note: You can find more information in the FC-9000 Redbook, at <http://www.ibm.com/redbooks>.

IOCP considerations

CNT (Inrange FC-9000) Switch ID Definition — No offsets on the switch ID or port address, but every value must be in hex for the mainframe.

Note: Domain ID 32 (the default) would be x'20'.

Security Appliances

This chapter contains information on the following security appliances:

- ◆ Overview 102
- ◆ Decru DataFort FC-Series security appliance 104
- ◆ Neoscale CryptoStore security appliance..... 118

Overview

This section contains the recommended settings and topology for the Decru DataFort FC-Series security appliance and the Neoscale Crypto security appliance that sit in the storage data path. Refer to the vendor's user guide and application guide for exact implementation steps. The following recommendations are valid only for local (proximate) Fibre Channel based fabrics with block transfer storage (for example, CLARiiON and Symmetrix). These appliances implement encryption/decryption based on SCSI FCP addressing (WWN, PID, LUN, LBA). Therefore, any changes of the initiator WWN, remapping of the LUN ID, or data movement from one partition to another requires user intervention for keys management. Key Management in each Decru appliance is implemented using the Lifetime Key Management Server. For more information about key management (duplication, restore, transport, etc.), refer to the vendor's user or system administration guide.

Data mobility in the encrypted form requires data consistency. Applications must commit data transactions through any of these appliances prior to data movement. For example, a volume copy should be prepared for backup activity by coalescing the database, un-mounting or synchronizing the file system, and shutting down or quiescent relevant applications.

Disk encryption mechanism

The Advanced Encryption Standard (AES) or Data Encryption Standard (DES) encryption algorithm of the data-at-rest in the disk is restricted on the Open System Logical Block Address (LBA) boundary, which is 512 bytes block. Both Decru and Neoscale do not add encryption rule (metadata) into the encrypted block, hence there is always a one-to-one correspondence from the unencrypted (clear text) block of data to the encrypted (crypto text) block of data.

Each vendor implements different algorithm to ensure that each of the 512 bytes block is encrypted using different keys in order to ensure the security of the data. The algorithm involves the position of the data block in the disk that is based on the LBA. As the result, the host can randomly access any LBA location through the disk encryption appliance to randomly retrieve the unencrypted data from the encrypted disk.

Impact of the storage application to the encrypted data

The disk encryption appliance provides the rule of how a particular key applies to a set of the data. This rule affects the SCSI addressing level, such as the I-T-L (Initiator-Target-LUN Nexus). The encrypted dataset does not carry any information about the rule (metadata), therefore the key must be appropriately associated with the same dataset regardless of the I-T-L mapping.

The storage application data (such as Snapview, Mirrorview, Clone, and SRDF) provides a point-in-time duplicate of data, or a synchronous duplication of data. These applications guarantee that the data is duplicated or transferred as the whole block of the Logical Block Address. The encrypted data can be transferred into another volume with different association of the I-T-L. The user must be actively involved to ensure the correct association of the key with the dataset because of the loose correlation of the keys to the dataset.

Follow these rules for encrypted data mobility:

- ◆ The encrypted data must be transferred in the whole entirety of the dataset.
- ◆ The encrypted data must be transferred to the same location of the LBA vicinity from the original source

Impact of online encryption data preparation and online encryption data re-keying with the storage applications

Consider the following scenarios between disk encryption and storage application:

- ◆ The data is currently converted from clear text to the crypto text.
- ◆ The encrypted data is replaced with other encrypted data with different keys.

In the above scenarios, part of the dataset is in either clear text (unencrypted) or crypto text (encrypted) with different key. As a result, the integrity of the point-in-time snapshot or synchronous replication of this dataset *cannot* be guaranteed.

Additional information on storage security can be found on EMC's [Powerlink](#). Select best practices for implementing various secure SAN mechanisms are contained throughout this chapter.

Decru DataFort FC-Series security appliance

The information in this section contains the recommended settings and topology for the implementation of the security appliance that sits in the path of the storage. Refer to the vendor's user guide and application guide for the exact implementation steps.

Recommendations are only valid for local (proximate) Fibre Channel-based fabric and disk-based (block transfer) storage (CLARiiON and Symmetrix).

Both Decru appliances implement encryption/decryption based on FCP (SCSI over Fibre Channel) addressing (WWN, PID, LUN, LBA). Therefore, any changes of the initiator WWN, remapping of the LUN ID, or data movement from one partition to another, requires user intervention for keys management. For more information about keys management (duplication, restore, transport, etc.), refer to each vendor's user or system administration guide.

Data mobility in the encrypted form requires data consistency. Therefore, any application must commit the data transaction (write to disk) through any of these appliances prior to data movement. In other words, the volume copied should be prepared for backup activity (coalesce database, file-system unmounted or synchronized, application shutdown or quiescent).

Key management in each Decru appliance is implemented using the Lifetime Key Management Server.

This section provides the following information on the Decru DataFort FC-series security appliance:

- ◆ [“Decru virtualization” on page 105](#)
- ◆ [“Decru mapping for the encrypted storage” on page 107](#)
- ◆ [“Decru Cryptainers vault” on page 108](#)

Following this discussion, further information is provided on supported configurations and best practices for:

- ◆ [“Decru DataFort FC-52x – Disk” on page 109](#)
- ◆ [“Decru DataFort FC-52x – Tape” on page 113](#)

Decru virtualization

The Decru DataFort appliance operates as a FC proxy for either the tape or disk encryption. As a FC proxy, all the ports in the DataFort FC-520, FC-525, and FC-1020 (both the host and storage ports) will assume either an N_Port or an NL_Port. Both the storage ports can automatically negotiate between 1 Gb/s and 2 Gb/s. If the Decru appliance is configured as a FC proxy that provides a single FC node, then the ports are configured as N_Port (Decru's terminology is *Single-ID mode*). In this mode, the Decru can only emulate one node from both host and storage ports.

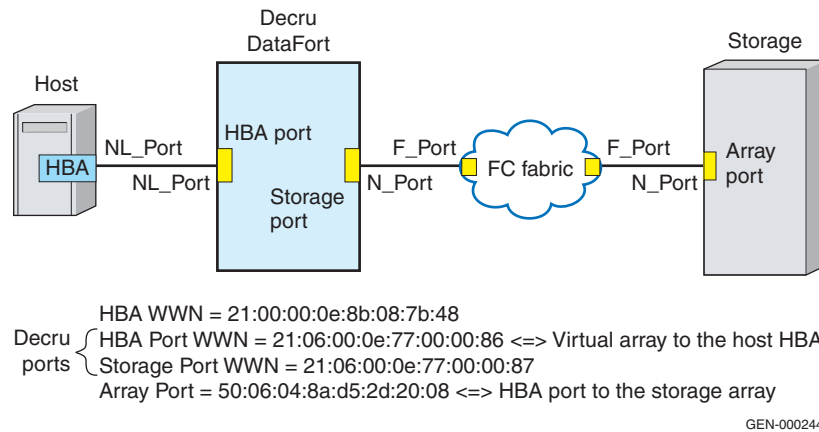
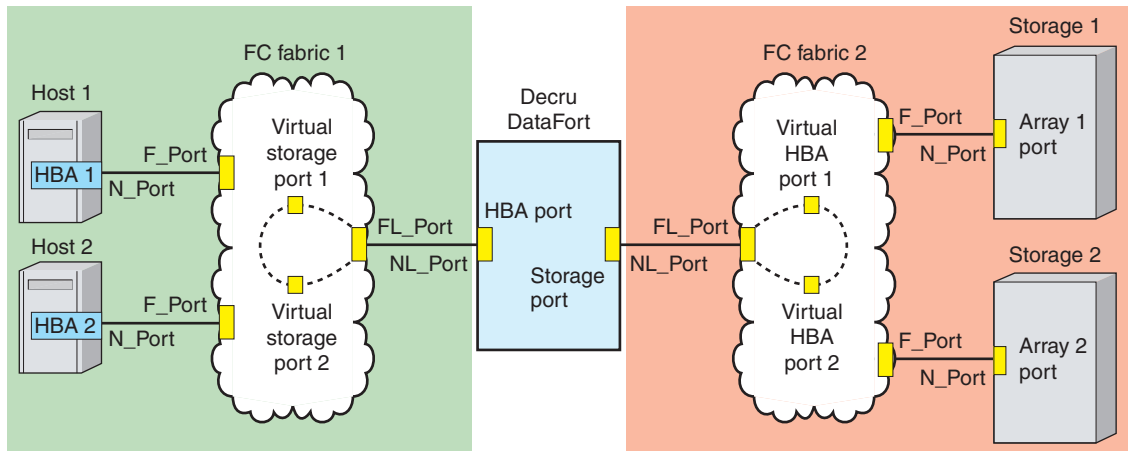


Figure 49 Single-ID mode (virtualization disabled)

If the Decru DataFort is configured as an FC proxy that provides emulation for multiple FC nodes, then the port will be configured as a NL_Port (Decru's terminology is *Multi-ID mode*). As the result of multi-ID mode, Decru DataFort will present several loop nodes that will emulate virtual nodes. If a storage port is virtualized, DataFort will then present a **Virtual Device** to the host through the host port. If a host port or an HBA port is virtualized, DataFort will then present a **Virtual Host** to the storage through the storage ports.



GEN-000245

HBA 1 WWN = 21:00:00:0e:8b:08:7b:48 <=> Virtual HBA port 1 WWN = 21:06:00:0e:77:00:01:87
 HBA 2 WWN = 21:00:00:0e:9b:08:7a:92 <=> Virtual HBA port 2 WWN = 21:06:00:0e:77:00:02:87
 DECRU HBA port WWN = 21:06:00:0e:77:00:00:86

Array port 1 WWN = 50:06:04:8a:d5:2d:20:08 <=> Virtual Storage port 1 = 21:06:00:0e:77:00:01:86
 Array port 2 WWN = 50:06:04:8a:d5:2e:10:08 <=> Virtual storage port 2 = 21:06:00:0e:77:00:02:86
 DECRU storage port WWN = 21:06:00:0e:77:00:00:87

Figure 50 Multi-ID mode (virtualization enabled)

Decru virtualization will always utilize a NL_Port to represent or virtualize multiple devices connected to its host or storage ports. To accomplish virtualization, DataFort uses the emulation of private loops in each of its ports. As a result, the FC switch must be able to translate the addresses of the emulated devices in the Arbitrated Loop to the address of the FC fabric. The typical switch port that translate between the arbitrated loop address and the fabric address is identified as an FL_Port. Hence, the Decru virtualization configuration requires the support of the FL_Port on the FC switch. For direct connection to either the host, storage, or tape ports, the Decru DataFort ports must be configured as a NL_Port (Decru's terminology is *Multi-ID mode*). The HBA port, storage, or tape ports must be configured as L_Port.

For complete information about configuring the virtualization settings, please see the "DataFort Setting and Status" section in the *Decru Fibre Channel Series Administration Guide*.

Decru mapping for the encrypted storage

The concept of virtualization or emulation of a storage node in the storage port allows the mapping of a WWN storage node to a WWN of a virtual node in the DataFort host port (Decru's terminology is *Port-Mapping*). Port-mapping preserves the same LUN numbering from the storage node to the virtualized storage node in the DataFort. This feature is very useful in the disk environment where multipathing software or management software is being used.

For tapes, it is useful to be able to sync up backup apps, etc, to consistent LUN numbering. However, port-mapping is less important in the tape environment unless there are multiple tape libraries that each contain both the robot LUN and multiple tape LUNs.

Please note that the number of ports that can be mapped depends on the DataFort capacity for mapping. For more information on the capacity of the Decru DataFort please see the "Planning the Network Configuration" section in the *Decru Fibre Channel Series Administration Guide*.

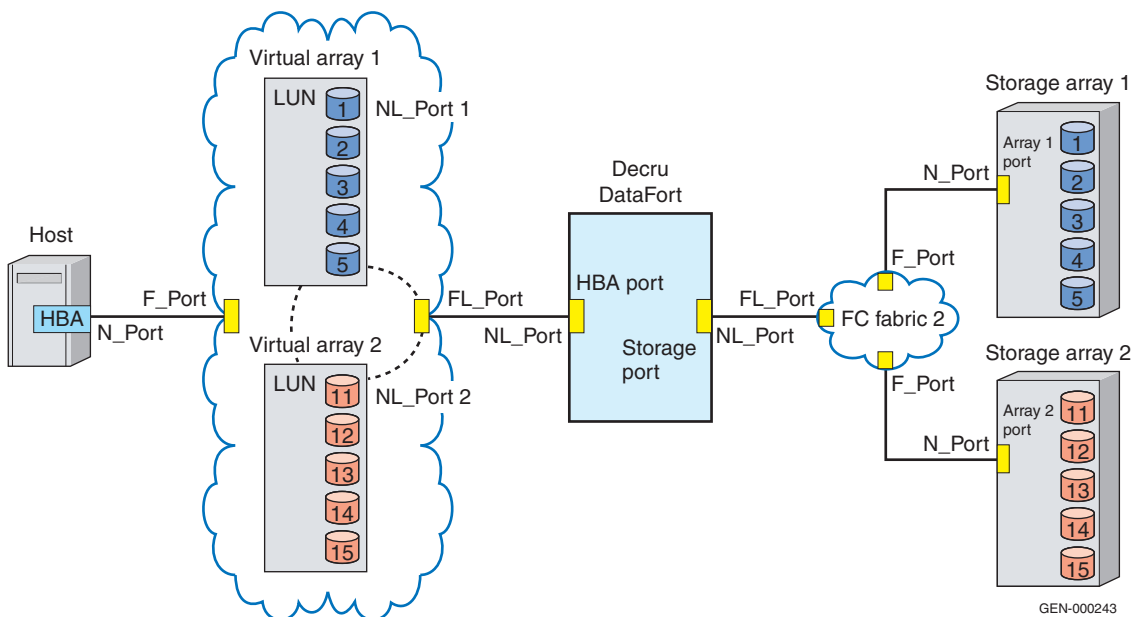


Figure 51 DataFort port mapping (storage side virtualization enabled)

Another option of DataFort mapping is known as *LUN-mapping*. LUN-mapping allows DataFort to assign a virtual LUN number to each device port (such as, storage port, robotic arm port, or tape drive port). As a result of LUN-mapping, DataFort can present a virtual storage node with multiple LUNs. For example, LUN0 can represent a robot arm while LUN1, LUN2, and LUN3 can represent tape drives. In reality, each of these LUNs are actually FC ports that are part of the tape library.

Another use of LUN-mapping is to provide the capability to map more storage ports than the maximum number of virtual devices that DataFort can virtualize.

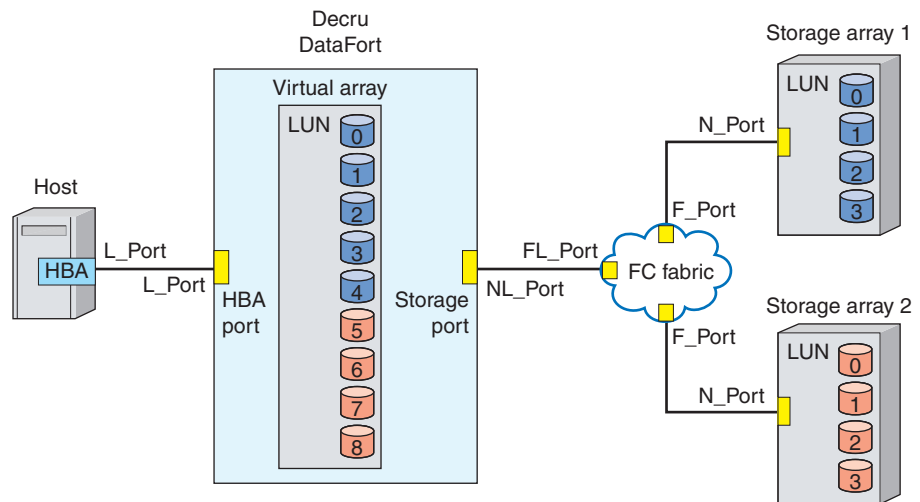


Figure 52 DataFort LUN mapping (storage side virtualization enabled)

Decru Cryptainers vault

Decru DataFort uses a concept called *Cryptainers* in order to manage key assignments, access control, and other security-related features. The Cryptainers are associated with storage WWNs and LUNs. Information that is related to Cryptainers is stored in a specific region of encrypted storage that has default properties in a tape environment and configurable properties in disk storage environments. The Decru DataFort term for the specific region is *Metadata*. DataFort appliances automatically associates the

Cryptainers to the storage LUN when the Metadata exist. Please refer to the "Introduction of Decru" section in the *Decru Fibre Channel Series Administration Guide* for further description of the Cryptainers.

Decru DataFort FC-52x – Disk

Supported configurations

The supported DataFort configuration for EMC disk arrays implements *storage virtualization* and *host virtualization* mode. DataFort architecture allows up to seven hosts and seven storage devices to be virtualized by one DataFort.

Storage virtualization is the act of virtualizing the storage in the DataFort host-side (storage) port. This is the default setting on the DataFort. The storage virtualization mode means that hosts actually see multiple Arbitrated Loop Nodes. Each of these Arbitrated Loop Nodes corresponds to a virtualized EMC storage port. In order to achieve this configuration, the virtualization for the storage and host port must be set to ON. The use of the storage virtualization mode enables DataFort to encrypt EMC storage using the port-mapped Cryptainer (one Cryptainer vault per LUN). With a Symmetrix array, LUN 0, which does not need to be encrypted, is used as the gateway for the Symmetrix masking application.

Host virtualization is the act of virtualizing a host on the DataFort storage-side [initiator] port. Host virtualization is *not* enabled by default. Enabling host virtualization is recommended especially for clustered hosts or shared volume environments having ≤ 7 hosts. Note that virtualization is only supported on FC switches that support FL connections or on storage ports that support AL connections. Please refer to Decru DataFort Disk best practices and release notes for more information.

Host virtualization is beneficial, or required, for the following common scenarios:

- ◆ Multiple hosts sharing same LUN.
- ◆ Preserving existing storage array masking by combining host virtualization with the option to forward host's WWN.

Note: Using the WWN forwarding feature is supported only if hosts and storage ports are on independent fabrics or VSANs, or if hard zoning is used.

- ◆ LUN masking in storage arrays based on host OS type.

Host virtualization allows you to specify different host types for each virtualized host presented on the DataFort initiator loop, which is useful for this scenario.

Please refer to the *Decru Fibre Channel Series Administration Guide* for details regarding configuration of virtualization and mapping.

Note: Migration to an encrypted storage environment utilizing Decru DataFort is disruptive in nature. The disruption occurs due to zoning requirements for host-to-virtual storage nodes, and virtual host nodes-to-storage ports. In addition, the migration also requires changing the access control setting of the storage, to the Decru DataFort storage port. After zoning changes, cryptainer assignment, and granting of access controls, Decru DataFort for Disk does not affect any host applications.

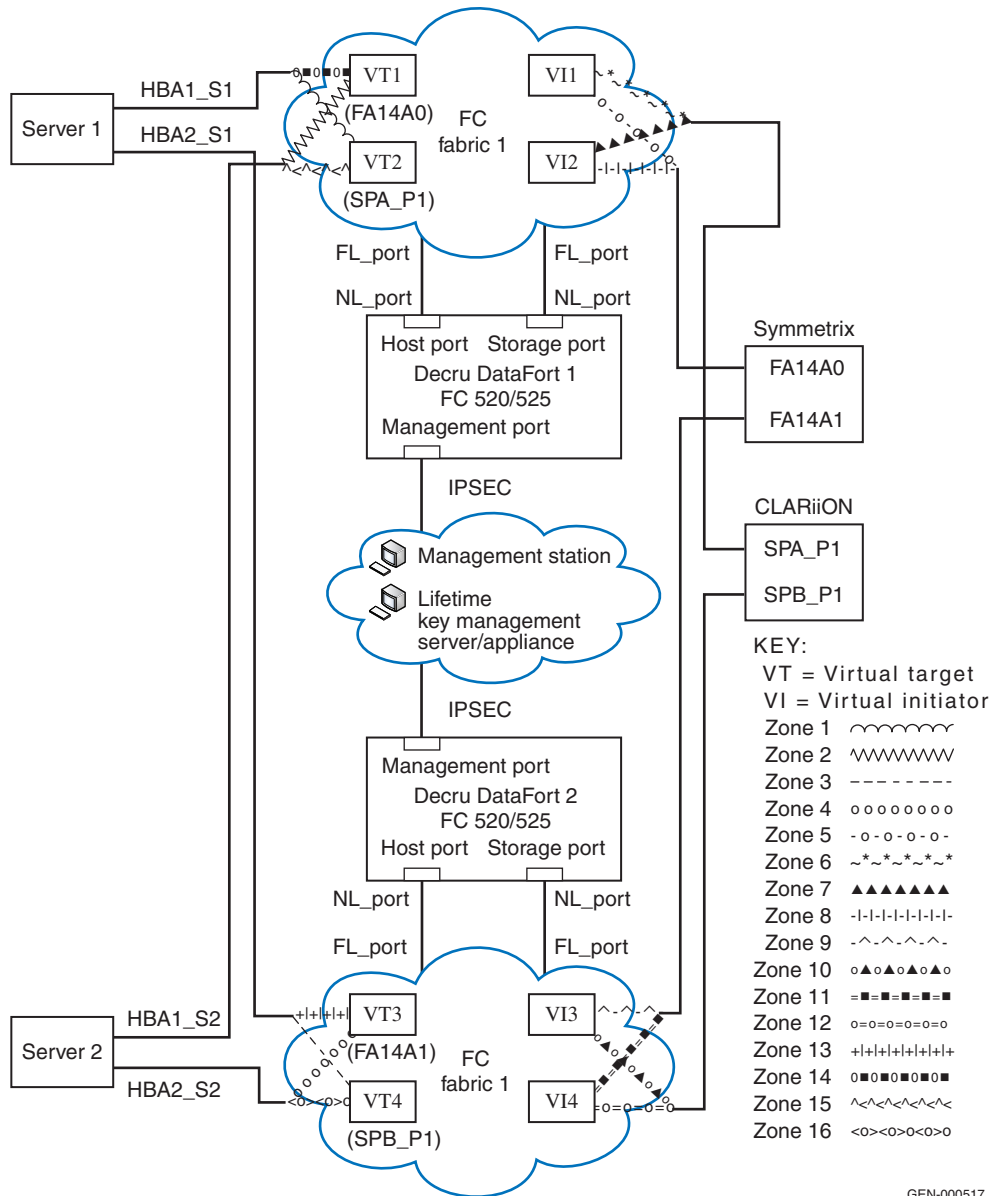
For complete information about configuring Decru DataFort with EMC storage, refer to the "Troubleshooting" section in the *Decru Fibre Channel Series Administration Guide*.

In order to provide a high availability disk encryption configuration, a cluster of minimally two DataForts must be implemented in a high availability SAN configuration. Refer to the "DataFort Initialization" section in the *Decru Fibre Channel Series Administration Guide* for more information about setting up the cluster.

Based on a clustered DataFort configuration, the following list contains the supported EMC storage configurations:

- ◆ EMC Symmetrix host bit settings must be set to default Symmetrix Fibre Channel open system host bit settings (VCM, PtoP, auto speed, and SC3 enabled).
- ◆ CLARiiON host settings must be similar to the default CLARiiON Fibre Channel open system host settings (host type 1, failover mode 1, communication port enabled).
- ◆ EMC Symmetrix with multipathing (at least one EMC Symmetrix port per path).
- ◆ EMC CLARiiON with multipathing (at least one corresponding SP port per path).
- ◆ EMC PowerPath is required for high availability configurations for both Symmetrix and CLARiiON. Host multi-pathing software is required to route the IO between paths because clustered DataForts do not route IO between multipaths.

- ◆ CLARiiON requires PowerPath to support nondisruptive upgrades (NDU).



GEN-000517

Figure 53 Recommended configuration of Decru DataFort with EMC storage

The following are recommended SAN topologies:

- ◆ Fibre Channel switches that support NL_Port nodes
- ◆ Brocade Silkstorm
- ◆ Cisco MDS
- ◆ Brocade M Series non-director switches, such as Sphereon 4500

When using DataFort with a Brocade M Series director, implement an edge switch to provide the FL_Port connectivity.

Note: All switch ports require auto-topology or FL_Port settings.

Best practices

This section lists best practices to consider.

- ◆ Although zone configuration can use both WWN zoning and port zoning, best practice suggests port zoning that allows multiple loop IDs in a single port.
- ◆ DataFort must be in storage virtualization mode, where the host is actually seeing multiple loop IDs, where each loop ID corresponds to a virtualized EMC storage port.
- ◆ Decru recommends adding all possible AL_PAs that can exist on the host port, such as WWN zoning (six different WWNs) for the port connected to the host. For example:

```
NL          a90725;          3;21:06:00:0e:77:00:00:86;20:06:00:0e:77:00:00:86;na
PC4s: PCP [Decru          _Secure_          0001]
Fabric Port Name: 20:07:00:60:69:90:02:bd
NL          a90726;          3;21:04:00:0e:77:00:00:86;20:04:00:0e:77:00:00:86;na
PC4s: PCP [Decru          _Secure_          0001]
Fabric Port Name: 20:07:00:60:69:90:02:bd
NL          a90727;          3;21:02:00:0e:77:00:00:86;20:02:00:0e:77:00:00:86;na
PC4s: PCP [Decru          _Secure_          0001]
Fabric Port Name: 20:07:00:60:69:90:02:bd
NL          a90729;          3;21:00:00:0e:77:00:00:86;20:00:00:0e:77:00:00:86;na
PC4s: PCP [Decru          _Secure_          0001]
Fabric Port Name: 20:07:00:60:69:90:02:bd
```

- ◆ Cryptainer mapping for EMC storage must use port mapping. As a result, each EMC storage port will be mapped to a virtual storage port. For high availability, map the LUNs from the secondary storage port as the alternate paths to the associated Cryptainers of the same LUN from the primary storage.

The *Decru DataFort Fibre Channel Series Administration Guide* contains more information about storage virtualization mode and cryptainer mapping.

- ◆ Cryptainers that provide security in the path between host and storage must be clustered together.
- ◆ DataFort provides a method to combine the references for any of the host and storage ports as a *Collection*. For complete information about configuring the Decru DataFort with EMC storage, please review the "Troubleshooting" section in the *Decru Fibre Channel Series Administration Guide*.
- ◆ Access control implementation with EMC storage:
 - Requires modification of the zone and the access control list of the Symmetrix and CLARiiON arrays.
 - Requires all EMC storage LUNs to grant the access from DataFort Storage ports.
 - Uses the DataFort GUI/CLI for individual host access to the secured LUNs controlled by the DataFort Cryptainers.
- ◆ Each virtualized storage port can support up to 255 LUNs (one Cryptainer per LUN).
- ◆ Be aware of the following configuration limits:
 - Seven storage ports per Decru box, because DataFort can virtualize only seven storage ports to the host.
 - Maximum fan-out for non-virtualized host ports to the DataFort port is 32 initiator ports (Fibre Channel HBA ports) per Decru box.

Note: The introduction of a high fan-out number will increase the latency of frame transmission. The user must be aware of the scalability impact to the host application, path recovery, and storage application.

References

- ◆ The *Decru Fibre Channel Series Administration Guide* contains more information about storage virtualization mode and cryptainer mapping.
- ◆ The *Decru DataFort Fibre Channel Series CLI Guide*.

Decru DataFort FC-52x – Tape

Decru DataFort FC-1020 provides five independent Fibre Channel host ports and five independent storage ports. However, the DataFort FC-1020 host and storage ports are engineered only for tape

Supported configuration

encryption (no disk encryption). The FC-520, FC525, and FC-1020 can be clustered together to provide high availability for tape encryption.

The supported DataFort configuration for EMC tape storage is the *Storage Virtualization Mode*. However, since tape is removable media, Decru DataFort provides a different scope of key assignments such as *Tape Pool Collection*, *Tape Pool*, or *Tape Device*. As a result, the Cryptainer configuration is different from the disk configuration. Refer to the "Tape Management" section of the *Decru Fibre Channel Series Administration Guide* for more information on tape key management. For supported configuration information, refer to the *Decru DataFort FC520, FC525 for Disk*.

Note: Migration to the encrypted storage environment utilizing the Decru DataFort is disruptive in nature. The disruption happens due to zoning requirements for host-to-virtual storage nodes and virtual host nodes-to-storage ports. In addition, the migration also requires the storage to change the access control setting to the Decru DataFort storage port (some virtual tape library environments). After the zoning changes, cryptainer assignment, and granting of access controls, Decru DataFort for tape will not affect any of the host applications. For the complete information about configuring Decru DataFort with EMC storage, refer to the "Troubleshooting" section in the *Decru Fibre Channel Series Administration Guide*.

An example of a topology without DataFort is shown in [Figure 54 on page 115](#).

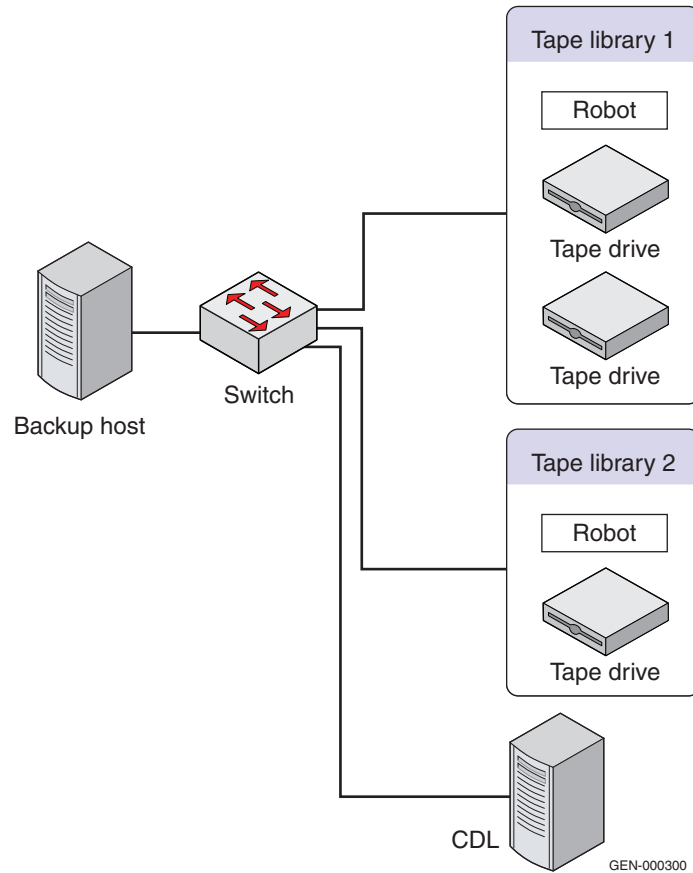


Figure 54 Topology without DataFort example

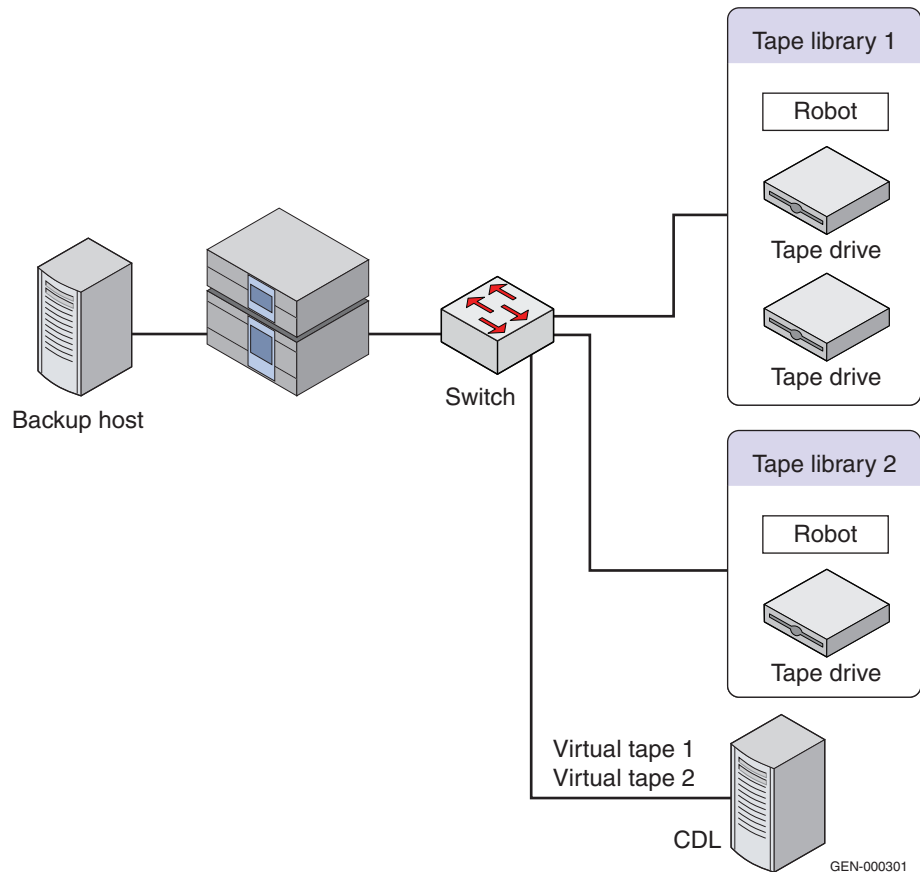


Figure 55 Virtualization enabled with port mapping

Best practices

This section lists best practices to consider.

- ◆ Although the DataFort is capable of providing separate paths for high availability configurations with tape encryption, DataFort depends on either user intervention or on an application mechanism to provide failover from one storage path to another. Therefore the high availability configuration for DataFort tape encryption should be recognized as hot-spare.
- ◆ Decru DataFort provides compression for the tape media in addition to encryption. Due to the nature of encryption, compression of cipher-text data can actually expand the data set.

Because of the proxy nature of the DataFort appliance, DataFort enables or disables the compression based on the host application setting.

- ◆ If the host HBA accommodates FC-Tape, Decru DataFort tape encryption appliances can accommodate the FC-Tape recovery method. It is recommended that you enable FC-Tape negotiation because the recovery for the disruption using FC-Tape is more efficient.

For further tips, limitations, and workarounds refer to the *Decru DataFort Fibre Channel Series Release Notes* and the "Troubleshooting" section of the *DecruFibre Channel Series Administrator Guide*.

References

- ◆ The *Decru Fibre Channel Series Administration Guide* contains more information about storage virtualization mode and cryptainer mapping.
- ◆ The *Decru Fibre Channel Series CLI Guide*.

Neoscale CryptoStore security appliance

This section contains the recommended settings and topology for the implementation of the security appliance that sits in the path of the storage. Refer to the vendor's user guide and application guide for the exact implementation steps. Recommendations are valid only for local (proximate) Fibre Channel-based fabric and disk-based (block transfer) storage (CLARiiON and Symmetrix).

Both of the Neoscale CryptoStore appliances implement encryption/decryption based on FCP (SCSI over Fibre Channel) addressing (WWN, PID, LUN, LBA). Therefore, any changes of the initiator WWN, remapping of the LUN ID, or data movement from one partition to another, requires user intervention for keys management. For more information about keys management (duplication, restore, transport, etc.), refer to each vendor's user or system administration guide.

Data mobility in the encrypted form requires data consistency. Therefore, any application must commit the data transaction (write to disk) through any of these appliances prior to data movement. In other words, the volume copied should be prepared for backup activity (coalesce database, file-system un-mounted or synchronized, application shutdown or quiescent).

This section contains security implementation information for the following:

- ◆ [“Neoscale CryptoStor FC-2002 for Disk,”](#) next
- ◆ [“Neoscale CryptoStor FC702/704 for Tape”](#) on page 121

Neoscale CryptoStor FC-2002 for Disk

CryptoStor FC2002 is a security appliance that provides *FC pass-through* behavior. *FC pass-through* has the following characteristics:

- ◆ The FC2002 appliance does not provide ports that model any of the traditional Fibre Channel topologies such as: N_Port, L_Port, F_Port, E_Port, B_Port, and so on.

- ◆ The FC2002 ports provide speed auto-negotiation with the connected port. The negotiated speed is the common highest speed for both of the ports. The FC2002 is not capable of providing different speeds between the two ports.
- ◆ FC 2002 ports carry the mode of the port-pairs without introducing any Fibre Channel topology termination. Hence, connecting both of the ports to the same fabric may introduce unknown behavior to the related switch.

Note: Migration to the encrypted disk environment for Neoscale CryptoStor FC-2002 *does not* cause lengthy disruption. Because of the transparent nature of the FC2002, the disruption is limited to the cabling of the CryptoStor FC2002 appliance between the switch port and the storage port. For complete information about configuring the Neoscale CryptoStor FC2002W, refer to the "Overview" section in the *CryptoStor Fibre Channel Series Administration Guide*.

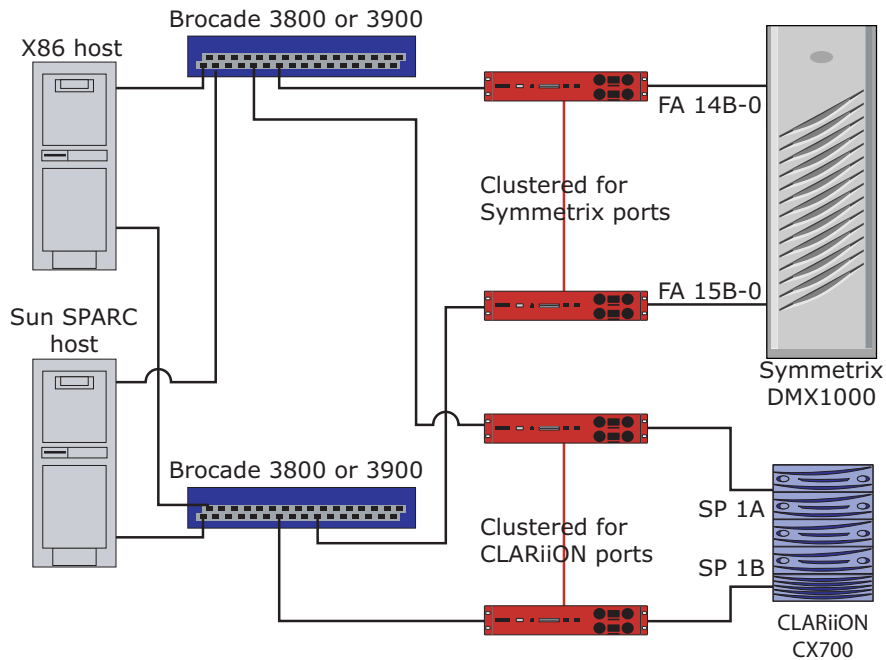
In order to provide a high-availability disk encryption configuration, a cluster of a minimum two CryptoStor FC2002 appliances must be implemented in the High Availability SAN configuration. Refer to the "System Administrator Handbook" section of the *CryptoStor Fibre Channel Series Administration Guide* for more information about setting up the cluster. The following list contains EMC-supported storage configurations based on clustered CryptoStor configurations:

- ◆ EMC Symmetrix configuration bit according to recommendation for the platform operating system of the connected host. Please refer to "Symmetrix Fibre Bit Settings" in the [EMC Support Matrix](#).
- ◆ EMC CLARiiON configuration setting according to the recommendation for the operating system platform of the connected host. Refer to host connectivity guides for each operating system platform.
- ◆ EMC Symmetrix with multipathing (at least one EMC Symmetrix port per path).
- ◆ EMC CLARiiON with multipathing (at least one corresponding SP port per path).
- ◆ EMC PowerPath required for high availability configurations for both Symmetrix and CLARiiON. The host multipathing software is required to route the IO between the paths.

- ◆ CLARiiON requires PowerPath to support non-disruptive upgrades (NDU).

Same encryption keys must be applied to the redundant Symmetrix and CLARiiON ports.

Encrypted Symmetrix LUNs must be mapped to both FA 14B-0 and FA 15B-0.



Paths between arrays and FC-2002s are clear text. All other paths are encrypted text.

Figure 56 Recommended CryptoStor 2002 configuration for EMC storage products

Supported configurations

The following are supported SAN configurations:

- ◆ Neoscale FC2002 connected inline between the hosts and the fabric (N_Port or NL_Port)
- ◆ Neoscale FC2002 connected inline between the fabric and the storage (N_Port or NL_Port)

Note: CryptoStor FC-2002 is a true transparent appliance. Therefore, there is no way to zone the CryptoStor FC-2002 ports.

Note: Zone changes and volume access control changes are not necessary between the existing host and the existing storage port.

Limitations Note the following configuration limits:

- ◆ The limit of the host fan-out depends on the EMC storage capability.

Note: The introduction of a high fan-out number will increase the latency of frame transmission. The user must be aware of the scalability impact to the host application, path recovery, and storage application.

- ◆ All possible paths where the host port is connected to the storage port must be secured by a clustered CryptoStor FC-2002.
- ◆ Each LUN in the redundant paths must use the same Data Encryption Policy (therefore, the same key).

Best practices Clustered Cryptostor FC 2002 must be connected on every single path where the host IO can be routed.

- References**
- ◆ The *CryptoStor Fibre Channel Series Administration Guide* contains detailed information on setting the data encryption policy.
 - ◆ *CryptoStor Fibre Channel Series Technical Reference Manual*.

Neoscale CryptoStor FC702/704 for Tape

Neoscale CryptoStor for Tape functions as the proxy device between the backup host and the tape library. As the proxy device, both the CryptoStor FC702/FC704 host, and storage ports can be configured as N_Port (Point-to-Point), NL_Port (Arbitrated Loop), and auto (try Loop first before Point-to-Point). In addition to the topology, the host and storage ports are capable of supporting 1 GB/s and 2 GB/s Fibre Channel speeds. By default all of the ports are auto topology and auto speed negotiation. The tape port virtualization requires that the host port must be configured as the NL_Port in order to emulate the tape library.

Supported configuration

The following are supported SAN configurations:

- ◆ The CryptoStor FC702/FC704 supported configuration is multi-ID on the host port (NL_Port), and single-ID on the target port (N_Port). Using this configuration, multiple tape libraries can be emulated as a single virtual device.




Figure 57 CryptoStor Tape 700 deployment example

Note: Migration to the encrypted tape environment for Neoscale CryptoStor FC Tape 700 series is disruptive in nature. The disruption happens due to the requirement of the change for the zoning both for the hosts to the virtual storage nodes and the virtual host node to the tape library ports. Furthermore, the migration also requires the tape library to change the access control setting to the Neoscale CryptoStor FC 700 series storage port (some virtual tape library environments). After the zoning changes, volume pools assignment, rules assignment, and access controls had been granted, Neoscale CryptoStor for Tapes will not affect any of the host applications. For complete information about configuring the Neoscale CryptoStor Tape 7000 series, please review the "Overview" section in the *DecruFibre Channel Series Administrator Guide*.

Note: CryptoStor FC-2002 is a true transparent appliance. Therefore, there is no way to zone the CryptoStor FC-2002 ports.

Note: Zone changes and volume access control changes are not necessary between the existing host and the existing storage port.

CryptoStor FC tape encryption appliance use three level key hierarchies:

- ◆ System key — The highest level. The key resides within the appliance and can be exported only via Smartcard. It encrypts any keys and is stored on an internal disk, Secure NVRAM, or Active Memory.

- ◆ Encryption and Pool keys — Second level. Pool/Encryption keys are used to encrypt the tape keys. The encryption/pool keys are encrypted by the System key and are backed up in a configuration file as well as the key catalog. Pool keys can be exported as part of the configuration onto a smartcard or via a file in the catalog form.
- ◆ Tape keys — Third level. The tape keys are encrypted by the encryption/pool keys. The tape keys are used to encrypt data before writing it to tape. They can reside in encrypted form on either the media (cassette) or in the catalog.

Limitations Note the following configuration limits:

- ◆ The limit of the host fan-out depends on the EMC storage capability.

Note: The introduction of a high fan-out number will increase the latency of frame transmission. The user must be aware of the scalability impact to the host application, path recovery, and storage application.

- ◆ All possible paths where the host port is connected to the storage port must be secured by a clustered CryptoStor FC-2002.
- ◆ Each LUN in the redundant paths must use the same data encryption policy (therefore, the same key).

Best practices This section lists best practices to consider.

- ◆ The storage port can only support FC-tape if, and only if, the port is configured as point-to-point (Single ID). Therefore, host port must be configured as point-to-point when FC-tape is expected to be implemented against the tape library port.
- ◆ It is always considered a best practice to store the tape keys in the catalog away from the data they protect. However, since tape is a removable medium and due to the limited internal disk capacity of CryptoStor FC tape appliance, the customer might opt to store the tape keys on the tape itself. If the user opts to save the tape keys on the internal catalog of the CryptoStor, then a daily backup of the catalog is highly recommended. This, however, might very soon lead to lack of free space on the internal disk especially if the customer is using dynamic keys.

- References**
- ◆ *CryptoStor Tape FC 700 Administration Guide*
 - ◆ *CryptoStor Tape FC 700 Installation Guide*.

This glossary contains terms related to EMC products and EMC networked storage concepts.

A

access control

A service that allows or prohibits access to a resource. Storage management products implement access control to allow or prohibit specific users. Storage platform products implement access control, often called LUN Masking, to allow or prohibit access to volumes by Initiators (HBAs). *See also* “[persistent binding](#)” and “[zoning](#).”

active domain ID

The domain ID actively being used by a switch. It is assigned to a switch by the principal switch.

active zone set

The active zone set is the zone set definition currently in effect and enforced by the fabric or other entity (for example, the name server). Only one zone set at a time can be active.

agent

An autonomous agent is a system situated within (and is part of) an environment that senses that environment, and acts on it over time in pursuit of its own agenda. Storage management software centralizes the control and monitoring of highly distributed storage infrastructure. The centralizing part of the software management system can depend on agents that are installed on the distributed parts of the infrastructure. For example, an agent (software component) can be installed on each of the hosts (servers) in an environment to allow the centralizing software to control and monitor the hosts.

alarm	An SNMP message notifying an operator of a network problem.
any-to-any port connectivity	A characteristic of a Fibre Channel switch that allows any port on the switch to communicate with any other port on the same switch.
application	Application software is a defined subclass of computer software that employs the capabilities of a computer directly to a task that users want to perform. This is in contrast to system software that participates with integration of various capabilities of a computer, and typically does not directly apply these capabilities to performing tasks that benefit users. The term application refers to both the application software and its implementation which often refers to the use of an information processing system. (For example, a payroll application, an airline reservation application, or a network application.) Typically an application is installed "on top of" an operating system like Windows or Linux, and contains a user interface.
application-specific integrated circuit (ASIC)	A circuit designed for a specific purpose, such as implementing lower-layer Fibre Channel protocols (FC-1 and FC-0). ASICs contrast with general-purpose devices such as memory chips or microprocessors, which can be used in many different applications.
arbitration	The process of selecting one respondent from a collection of several candidates that request service concurrently.
ASIC family	Different switch hardware platforms that utilize the same port ASIC can be grouped into collections known as an ASIC family. For example, the Fuji ASIC family which consists of the ED-64M and ED-140M run different microprocessors, but both utilize the same port ASIC to provide Fibre Channel connectivity, and are therefore in the same ASIC family. For inter operability concerns, it is useful to understand to which ASIC family a switch belongs.
ASCII	ASCII (American Standard Code for Information Interchange), generally pronounced [æski], is a character encoding based on the English alphabet. ASCII codes represent text in computers, communications equipment, and other devices that work with text. Most modern character encodings, which support many more characters, have a historical basis in ASCII.
audit log	A log containing summaries of actions taken by a Connectrix Management software user that creates an audit trail of changes. Adding, modifying, or deleting user or product administration

values, creates a record in the audit log that includes the date and time.

authentication Verification of the identity of a process or person.

B

backpressure The effect on the environment leading up to the point of restriction. See [“congestion.”](#)

BB_Credit See [“buffer-to-buffer credit.”](#)

beaconing Repeated transmission of a beacon light and message until an error is corrected or bypassed. Typically used by a piece of equipment when an individual Field Replaceable Unit (FRU) needs replacement. Beaconing helps the field engineer locate the specific defective component. Some equipment management software systems such as Connectrix Manager offer beaconing capability.

BER See [“bit error rate.”](#)

bidirectional In Fibre Channel, the capability to simultaneously communicate at maximum speeds in both directions over a link.

bit error rate Ratio of received bits that contain errors to total of all bits transmitted.

blade server A consolidation of independent servers and switch technology in the same chassis.

blocked port Devices communicating with a blocked port are prevented from logging in to the Fibre Channel switch containing the port or communicating with other devices attached to the switch. A blocked port continuously transmits the off-line sequence (OLS).

bridge A device that provides a translation service between two network segments utilizing different communication protocols. EMC supports and sells bridges that convert iSCSI storage commands from a NIC-attached server to Fibre Channel commands for a storage platform.

broadcast Sends a transmission to all ports in a network. Typically used in IP networks. Not typically used in Fibre Channel networks.

- broadcast frames** Data packet, also known as a broadcast packet, whose destination address specifies all computers on a network. *See also "multicast."*
- buffer** Storage area for data in transit. Buffers compensate for differences in link speeds and link congestion between devices.
- buffer-to-buffer credit** The number of receive buffers allocated by a receiving FC_Port to a transmitting FC_Port. The value is negotiated between Fibre Channel ports during link initialization. Each time a port transmits a frame it decrements this credit value. Each time a port receives an R_Rdy frame it increments this credit value. If the credit value is decremented to zero, the transmitter stops sending any new frames until the receiver has transmitted an R_Rdy frame. Buffer-to-buffer credit is particularly important in SRDF and Mirror View distance extension solutions.
- C**
- Call Home** A product feature that allows the Connectrix service processor to automatically dial out to a support center and report system problems. The support center server accepts calls from the Connectrix service processor, logs reported events, and can notify one or more support center representatives. Telephone numbers and other information are configured through the Windows NT dial-up networking application. The Call Home function can be enabled and disabled through the Connectrix Product Manager.
- channel** With Open Systems, a channel is a point-to-point link that transports data from one point to another on the communication path, typically with high throughput and low latency that is generally required by storage systems. With Mainframe environments, a channel refers to the server-side of the server-storage communication path, analogous to the HBA in Open Systems.
- Class 2 Fibre Channel class of service** In Class 2 service, the fabric and destination N_Ports provide connectionless service with notification of delivery or nondelivery between the two N_Ports. Historically Class 2 service is not widely used in Fibre Channel system.
- Class 3 Fibre Channel class of service** Class 3 service provides a connectionless service without notification of delivery between N_Ports. (This is also known as datagram service.) The transmission and routing of Class 3 frames is the same

as for Class 2 frames. Class 3 is the dominant class of communication used in Fibre Channel for moving data between servers and storage and may be referred to as “Ship and pray.”

Class F Fibre Channel class of service	Class F service is used for all switch-to-switch communication in a multiswitch fabric environment. It is nearly identical to class 2 from a flow control point of view.
community	A relationship between an SNMP agent and a set of SNMP managers that defines authentication, access control, and proxy characteristics.
community name	A name that represents an SNMP community that the agent software recognizes as a valid source for SNMP requests. An SNMP management program that sends an SNMP request to an agent program must identify the request with a community name that the agent recognizes or the agent discards the message as an authentication failure. The agent counts these failures and reports the count to the manager program upon request, or sends an authentication failure trap message to the manager program.
community profile	Information that specifies which management objects are available to what management domain or SNMP community name.
congestion	Occurs at the point of restriction. See “backpressure.”
connectionless	Non dedicated link. Typically used to describe a link between nodes that allows the switch to forward Class 2 or Class 3 frames as resources (ports) allow. <i>Contrast with</i> the dedicated bandwidth that is required in a Class 1 Fibre Channel Service point-to-point link.
Connectivity Unit	A hardware component that contains hardware (and possibly software) that provides Fibre Channel connectivity across a fabric. Connectrix switches are example of Connectivity Units. This is a term popularized by the Fibre Alliance MIB, sometimes abbreviated to connunit.
Connectrix management software	The software application that implements the management user interface for all managed Fibre Channel products, typically the Connectrix -M product line. Connectrix Management software is a client/server application with the server running on the Connectrix service processor, and clients running remotely or on the service processor.

Connectrix service processor	An optional 1U server shipped with the Connectrix -M product line to run the Connectrix Management server software and EMC remote support application software.
Control Unit	In mainframe environments, a Control Unit controls access to storage. It is analogous to a Target in Open Systems environments.
core switch	Occupies central locations within the interconnections of a fabric. Generally provides the primary data paths across the fabric and the direct connections to storage devices. Connectrix directors are typically installed as core switches, but may be located anywhere in the fabric.
credit	A numeric value that relates to the number of available BB_Credits on a Fibre Channel port. <i>See</i> "buffer-to-buffer credit".
D	
DASD	Direct Access Storage Device.
default	Pertaining to an attribute, value, or option that is assumed when none is explicitly specified.
default zone	A zone containing all attached devices that are not members of any active zone. Typically the default zone is disabled in a Connectrix M environment which prevents newly installed servers and storage from communicating until they have been provisioned.
Dense Wavelength Division Multiplexing (DWDM)	A process that carries different data channels at different wavelengths over one pair of fiber optic links. A conventional fiber-optic system carries only one channel over a single wavelength traveling through a single fiber.
destination ID	A field in a Fibre Channel header that specifies the destination address for a frame. The Fibre Channel header also contains a Source ID (SID). The FCID for a port contains both the SID and the DID.
device	A piece of equipment, such as a server, switch or storage system.
dialog box	A user interface element of a software product typically implemented as a pop-up window containing informational messages and fields for modification. Facilitates a dialog between the user and the application. Dialog box is often used interchangeably with window.

DID An acronym used to refer to either Domain ID or Destination ID. This ambiguity can create confusion. As a result E-Lab recommends this acronym be used to apply to Domain ID. Destination ID can be abbreviated to FCID.

director An enterprise-class Fibre Channel switch, such as the Connectrix ED-140M, MDS 9509, or ED-48000B. Directors deliver high availability, failure ride-through, and repair under power to insure maximum uptime for business critical applications. Major assemblies, such as power supplies, fan modules, switch controller cards, switching elements, and port modules, are all hot-swappable.

The term director may also refer to a board-level module in the Symmetrix that provides the interface between host channels (through an associated adapter module in the Symmetrix) and Symmetrix disk devices. (This description is presented here only to clarify a term used in other EMC documents.)

DNS See “[domain name service name](#).”

domain ID A byte-wide field in the three byte Fibre Channel address that uniquely identifies a switch in a fabric. The three fields in a FCID are domain, area, and port. A distinct Domain ID is requested from the principal switch. The principal switch allocates one Domain ID to each switch in the fabric. A user may be able to set a Preferred ID which can be requested of the Principal switch, or set an Insistent Domain ID. If two switches insist on the same DID one or both switches will segment from the fabric.

domain name service name Host or node name for a system that is translated to an IP address through a name server. All DNS names have a host name component and, if fully qualified, a domain component, such as *host1.abcd.com*. In this example, *host1* is the host name.

dual-attached host A host that has two (or more) connections to a set of devices.

E

E_D_TOV A time-out period within which each data frame in a Fibre Channel sequence transmits. This avoids time-out errors at the destination Nx_Port. This function facilitates high speed recovery from dropped frames. Typically this value is 2 seconds.

E_Port	Expansion Port, a port type in a Fibre Channel switch that attaches to another E_Port on a second Fibre Channel switch forming an Interswitch Link (ISL). This link typically conforms to the FC-SW standards developed by the T11 committee, but might not support heterogeneous inter operability.
edge switch	Occupies the periphery of the fabric, generally providing the direct connections to host servers and management workstations. No two edge switches can be connected by interswitch links (ISLs). Connectrix departmental switches are typically installed as edge switches in a multiswitch fabric, but may be located anywhere in the fabric
Embedded Web Server	A management interface embedded on the switch's code that offers features similar to (but not as robust as) the Connectrix Manager and Product Manager.
error detect time out value	Defines the time the switch waits for an expected response before declaring an error condition. The error detect time out value (E_D_TOV) can be set within a range of two-tenths of a second to one second using the Connectrix switch Product Manager.
error message	An indication that an error has been detected. <i>See also</i> " information message " and " warning message ."
Ethernet	A baseband LAN that allows multiple station access to the transmission medium at will without prior coordination and which avoids or resolves contention.
event log	A record of significant events that have occurred on a Connectrix switch, such as FRU failures, degraded operation, and port problems.
expansionport	<i>See</i> " E_Port ."
explicit fabric login	In order to join a fabric, an Nport must login to the fabric (an operation referred to as an FLOGI). Typically this is an explicit operation performed by the Nport communicating with the F_port of the switch, and is called an explicit fabric login. Some legacy Fibre Channel ports do not perform explicit login, and switch vendors perform login for ports creating an implicit login. Typically logins are explicit.

F

- FA** Fibre Adapter, another name for a Symmetrix Fibre Channel director.
- F_Port** Fabric Port, a port type on a Fibre Channel switch. An F_Port attaches to an N_Port through a point-to-point full-duplex link connection. A G_Port automatically becomes an F_port or an E-Port depending on the port initialization process.
- fabric** One or more switching devices that interconnect Fibre Channel N_Ports, and route Fibre Channel frames based on destination IDs in the frame headers. A fabric provides discovery, path provisioning, and state change management services for a Fibre Channel environment.
- fabric element** Any active switch or director in the fabric.
- fabric login** Process used by N_Ports to establish their operating parameters including class of service, speed, and buffer-to-buffer credit value.
- fabric port** A port type (F_Port) on a Fibre Channel switch that attaches to an N_Port through a point-to-point full-duplex link connection. An N_Port is typically a host (HBA) or a storage device like Symmetrix or CLARiiON.
- fabric shortest path first (FSPF)** A routing algorithm implemented by Fibre Channel switches in a fabric. The algorithm seeks to minimize the number of hops traversed as a Fibre Channel frame travels from its source to its destination.
- fabric tree** A hierarchical list in Connectrix Manager of all fabrics currently known to the Connectrix service processor. The tree includes all members of the fabrics, listed by WWN or nickname.
- failover** The process of detecting a failure on an active Connectrix switch FRU and the automatic transition of functions to a backup FRU.
- fan-in/fan-out** Term used to describe the server:storage ratio, where a graphic representation of a 1:n (fan-in) or n:1 (fan-out) logical topology looks like a hand-held fan, with the wide end toward n. By convention fan-out refers to the number of server ports that share a single storage port. Fan-out consolidates a large number of server ports on a fewer number of storage ports. Fan-in refers to the number of storage ports that a single server port uses. Fan-in enlarges the storage capacity used by a server. A fan-in or fan-out rate is often referred to as just the

n part of the ratio; For example, a 16:1 fan-out is also called a fan-out rate of 16, in this case 16 server ports are sharing a single storage port.

FCP See ["Fibre Channel Protocol."](#)

FC-SW The Fibre Channel fabric standard. The standard is developed by the T11 organization whose documentation can be found at [T11.org](#). EMC actively participates in T11. T11 is a committee within the InterNational Committee for Information Technology (INCITS).

fiber optics The branch of optical technology concerned with the transmission of radiant power through fibers made of transparent materials such as glass, fused silica, and plastic.

Either a single discrete fiber or a non spatially aligned fiber bundle can be used for each information channel. Such fibers are often called optical fibers to differentiate them from fibers used in non-communication applications.

fibres A general term used to cover all physical media types supported by the Fibre Channel specification, such as optical fiber, twisted pair, and coaxial cable.

Fibre Channel The general name of an integrated set of ANSI standards that define new protocols for flexible information transfer. Logically, Fibre Channel is a high-performance serial data channel.

Fibre Channel Protocol A standard Fibre Channel FC-4 level protocol used to run SCSI over Fibre Channel.

Fibre Channel switch modules The embedded switch modules in the back plane of the blade server. See ["blade server"](#) on page 127.

firmware The program code (embedded software) that resides and executes on a connectivity device, such as a Connectrix switch, a Symmetrix Fibre Channel director, or a host bus adapter (HBA).

F_Port Fabric Port, a physical interface within the fabric. An F_Port attaches to an N_Port through a point-to-point full-duplex link connection.

frame A set of fields making up a unit of transmission. Each field is made of bytes. The typical Fibre Channel frame consists of fields: Start-of-frame, header, data-field, CRC, end-of-frame. The maximum frame size is 2148 bytes.

frame header	Control information placed before the data-field when encapsulating data for network transmission. The header provides the source and destination IDs of the frame.
FRU	Field-replaceable unit, a hardware component that can be replaced as an entire unit. The Connectrix switch Product Manager can display status for the FRUs installed in the unit.
FSPF	Fabric Shortest Path First, an algorithm used for routing traffic. This means that, between the source and destination, only the paths that have the least amount of physical hops will be used for frame delivery.
G	
gateway address	In TCP/IP, a device that connects two systems that use the same or different protocols.
gigabyte (GB)	A unit of measure for storage size, loosely one billion (10^9) bytes. One gigabyte actually equals 1,073,741,824 bytes.
G_Port	A port type on a Fibre Channel switch capable of acting either as an F_Port or an E_Port, depending on the port type at the other end of the link.
GUI	Graphical user interface.
H	
HBA	See "host bus adapter."
hexadecimal	Pertaining to a numbering system with base of 16; valid numbers use the digits 0 through 9 and characters A through F (which represent the numbers 10 through 15).
high availability	A performance feature characterized by hardware component redundancy and hot-swappability (enabling non-disruptive maintenance). High-availability systems maximize system uptime while providing superior reliability, availability, and serviceability.
hop	A hop refers to the number of InterSwitch Links (ISLs) a Fibre Channel frame must traverse to go from its source to its destination.

Good design practice encourages three hops or less to minimize congestion and performance management complexities.

host bus adapter A bus card in a host system that allows the host system to connect to the storage system. Typically the HBA communicates with the host over a PCI or PCI Express bus and has a single Fibre Channel link to the fabric. The HBA contains an embedded microprocessor with on board firmware, one or more ASICs, and a Small Form Factor Pluggable module (SFP) to connect to the Fibre Channel link.

I

I/O See [“input/output.”](#)

in-band management Transmission of monitoring and control functions over the Fibre Channel interface. You can also perform these functions out-of-band typically by use of the ethernet to manage Fibre Channel devices.

information message A message telling a user that a function is performing normally or has completed normally. User acknowledgement might or might not be required, depending on the message. See also [“error message”](#) and [“warning message.”](#)

input/output (1) Pertaining to a device whose parts can perform an input process and an output process at the same time. (2) Pertaining to a functional unit or channel involved in an input process, output process, or both (concurrently or not), and to the data involved in such a process. (3) Pertaining to input, output, or both.

interface (1) A shared boundary between two functional units, defined by functional characteristics, signal characteristics, or other characteristics as appropriate. The concept includes the specification of the connection of two devices having different functions. (2) Hardware, software, or both, that links systems, programs, or devices.

Internet Protocol See [“IP.”](#)

interoperability The ability to communicate, execute programs, or transfer data between various functional units over a network. Also refers to a Fibre Channel fabric that contains switches from more than one vendor.

- interswitch link (ISL)** Interswitch link, a physical E_Port connection between any two switches in a Fibre Channel fabric. An ISL forms a hop in a fabric.
- IP** Internet Protocol, the TCP/IP standard protocol that defines the datagram as the unit of information passed across an internet and provides the basis for connectionless, best-effort packet delivery service. IP includes the ICMP control and error message protocol as an integral part.
- IP address** A unique string of numbers that identifies a device on a network. The address consists of four groups (quadrants) of numbers delimited by periods. (This is called *dotted-decimal* notation.) All resources on the network must have an IP address. A valid IP address is in the form *nnn.nnn.nnn.nnn*, where each *nnn* is a decimal in the range 0 to 255.
- ISL** Interswitch link, a physical E_Port connection between any two switches in a Fibre Channel fabric.
- K**
- kilobyte (K)** A unit of measure for storage size, loosely one thousand bytes. One kilobyte actually equals 1,024 bytes.
- L**
- laser** A device that produces optical radiation using a population inversion to provide light amplification by stimulated emission of radiation and (generally) an optical resonant cavity to provide positive feedback. Laser radiation can be highly coherent temporally, spatially, or both.
- LED** Light-emitting diode.
- link** The physical connection between two devices on a switched fabric.
- link incident** A problem detected on a fiber-optic link; for example, loss of light, or invalid sequences.
- load balancing** The ability to distribute traffic over all network ports that are the same distance from the destination address by assigning different paths to different messages. Increases effective network bandwidth. EMC PowerPath software provides load-balancing services for server IO.

logical volume A named unit of storage consisting of a logically contiguous set of disk sectors.

Logical Unit Number (LUN) A number, assigned to a storage volume, that (in combination with the storage device node's World Wide Port Name (WWPN)) represents a unique identifier for a logical volume on a storage area network.

M

MAC address Media Access Control address, the hardware address of a device connected to a shared network.

managed product A hardware product that can be managed using the Connectrix Product Manager. For example, a Connectrix switch is a managed product.

management session Exists when a user logs in to the Connectrix Management software and successfully connects to the product server. The user must specify the network address of the product server at login time.

media The disk surface on which data is stored.

media access control *See* "MAC address."

megabyte (MB) A unit of measure for storage size, loosely one million (10^6) bytes. One megabyte actually equals 1,048,576 bytes.

MIB Management Information Base, a related set of objects (variables) containing information about a managed device and accessed through SNMP from a network management station.

multicast Multicast is used when multiple copies of data are to be sent to designated, multiple, destinations.

multiswitch fabric Fibre Channel fabric created by linking more than one switch or director together to allow communication. *See also* "ISL."

multiswitch linking Port-to-port connections between two switches.

N

name server (DNS) A service known as the distributed Name Server provided by a Fibre Channel fabric that provides device discovery, path provisioning, and

state change notification services to the N_Ports in the fabric. The service is implemented in a distributed fashion, for example, each switch in a fabric participates in providing the service. The service is addressed by the N_Ports through a Well Known Address.

- network address** A name or address that identifies a managed product, such as a Connectrix switch, or a Connectrix service processor on a TCP/IP network. The network address can be either an IP address in dotted decimal notation, or a Domain Name Service (DNS) name as administered on a customer network. All DNS names have a host name component and (if fully qualified) a domain component, such as *host1.emc.com*. In this example, *host1* is the host name and *EMC.com* is the domain component.
- nickname** A user-defined name representing a specific WWxN, typically used in a Connectrix -M management environment. The analog in the Connectrix -B and MDS environments is alias.
- node** The point at which one or more functional units connect to the network.
- N_Port** Node Port, a Fibre Channel port implemented by an end device (node) that can attach to an F_Port or directly to another N_Port through a point-to-point link connection. HBAs and storage systems implement N_Ports that connect to the fabric.
- NVRAM** Nonvolatile random access memory.
- O**
- offline sequence (OLS)** The OLS Primitive Sequence is transmitted to indicate that the FC_Port transmitting the Sequence is:
- a. initiating the Link Initialization Protocol
 - b. receiving and recognizing NOS
 - c. or entering the offline state
- OLS** See "[offline sequence \(OLS\)](#)".
- operating mode** Regulates what other types of switches can share a multiswitch fabric with the switch under consideration.

- operating system** Software that controls the execution of programs and that may provide such services as resource allocation, scheduling, input/output control, and data management. Although operating systems are predominantly software, partial hardware implementations are possible.
- optical cable** A fiber, multiple fibers, or a fiber bundle in a structure built to meet optical, mechanical, and environmental specifications.
- OS** *See "operating system."*
- out-of-band management** Transmission of monitoring/control functions outside of the Fibre Channel interface, typically over ethernet.
- oversubscription** The ratio of bandwidth required to bandwidth available. When all ports, associated pair-wise, in any random fashion, cannot sustain full duplex at full line-rate, the switch is oversubscribed.

P

- parameter** A characteristic element with a variable value that is given a constant value for a specified application. Also, a user-specified value for an item in a menu; a value that the system provides when a menu is interpreted; data passed between programs or procedures.
- password** (1) A value used in authentication or a value used to establish membership in a group having specific privileges. (2) A unique string of characters known to the computer system and to a user who must specify it to gain full or limited access to a system and to the information stored within it.
- path** In a network, any route between any two nodes.
- persistent binding** Use of server-level access control configuration information to persistently bind a server device name to a specific Fibre Channel storage volume or logical unit number, through a specific HBA and storage port WWN. The address of a persistently bound device does not shift if a storage target fails to recover during a power cycle. This function is the responsibility of the HBA device driver.
- port** (1) An access point for data entry or exit. (2) A receptacle on a device to which a cable for another device is attached.

port card	Field replaceable hardware component that provides the connection for fiber cables and performs specific device-dependent logic functions.
port name	A symbolic name that the user defines for a particular port through the Product Manager.
preferred domain ID	An ID configured by the fabric administrator. During the fabric build process a switch requests permission from the principal switch to use its preferred domain ID. The principal switch can deny this request by providing an alternate domain ID only if there is a conflict for the requested Domain ID. Typically a principal switch grants the non-principal switch its requested Preferred Domain ID.
principal switch	In a multiswitch fabric, the switch that allocates domain IDs to itself and to all other switches in the fabric. There is always one principal switch in a fabric. If a switch is not connected to any other switches, it acts as its own principal switch.
principle downstream ISL	The ISL to which each switch will forward frames originating from the principal switch.
principle ISL	The principal ISL is the ISL that frames destined to, or coming from, the principal switch in the fabric will use. An example is an RDI frame.
principle upstream ISL	The ISL to which each switch will forward frames destined for the principal switch. The principal switch does not have any upstream ISLs.
product	(1) Connectivity Product, a generic name for a switch, director, or any other Fibre Channel product. (2) Managed Product, a generic hardware product that can be managed by the Product Manager (a Connectrix switch is a managed product). Note distinction from the definition for “ device .”
Product Manager	A software component of Connectrix Manager software such as a Connectrix switch product manager, that implements the management user interface for a specific product. When a product instance is opened from the Connectrix Manager software products view, the corresponding product manager is invoked. The product manager is also known as an Element Manager.

product name A user configurable identifier assigned to a Managed Product. Typically, this name is stored on the product itself. For a Connectrix switch, the Product Name can also be accessed by an SNMP Manager as the System Name. The Product Name should align with the host name component of a Network Address.

products view The top-level display in the Connectrix Management software user interface that displays icons of Managed Products.

protocol (1) A set of semantic and syntactic rules that determines the behavior of functional units in achieving communication. (2) A specification for the format and relative timing of information exchanged between communicating parties.

R

R_A_TOV See "[resource allocation time out value.](#)"

remote access link The ability to communicate with a data processing facility through a remote data link.

remote notification The system can be programmed to notify remote sites of certain classes of events.

remote user workstation A workstation, such as a PC, using Connectrix Management software and Product Manager software that can access the Connectrix service processor over a LAN connection. A user at a remote workstation can perform all of the management and monitoring tasks available to a local user on the Connectrix service processor.

resource allocation time out value A value used to time-out operations that depend on a maximum time that an exchange can be delayed in a fabric and still be delivered. The resource allocation time-out value of (R_A_TOV) can be set within a range of two-tenths of a second to 120 seconds using the Connectrix switch product manager. The typical value is 10 seconds.

S

SAN See "[storage area network \(SAN\).](#)"

segmentation A non-connection between two switches. Numerous reasons exist for an operational ISL to segment, including interop mode incompatibility, zoning conflicts, and domain overlaps.

segmented E_Port	E_Port that has ceased to function as an E_Port within a multiswitch fabric due to an incompatibility between the fabrics that it joins.
service processor	See "Connectrix service processor."
session	See "management session."
single attached host	A host that only has a single connection to a set of devices.
small form factor pluggable (SFP)	An optical module implementing a shortwave or long wave optical transceiver.
SMTP	Simple Mail Transfer Protocol, a TCP/IP protocol that allows users to create, send, and receive text messages. SMTP protocols specify how messages are passed across a link from one system to another. They do not specify how the mail application accepts, presents or stores the mail.
SNMP	Simple Network Management Protocol, a TCP/IP protocol that generally uses the User Datagram Protocol (UDP) to exchange messages between a management information base (MIB) and a management client residing on a network.
storage area network (SAN)	A network linking servers or workstations to disk arrays, tape backup systems, and other devices, typically over Fibre Channel and consisting of multiple fabrics.
subnet mask	Used by a computer to determine whether another computer with which it needs to communicate is located on a local or remote network. The network mask depends upon the class of networks to which the computer is connecting. The mask indicates which digits to look at in a longer network address and allows the router to avoid handling the entire address. Subnet masking allows routers to move the packets more quickly. Typically, a subnet may represent all the machines at one geographic location, in one building, or on the same local area network.
switch priority	Value configured into each switch in a fabric that determines its relative likelihood of becoming the fabric's principal switch.

T

TCP/IP Transmission Control Protocol/Internet Protocol. TCP/IP refers to the protocols that are used on the Internet and most computer networks. TCP refers to the Transport layer that provides flow control and connection services. IP refers to the Internet Protocol level where addressing and routing are implemented.

toggle To change the state of a feature/function that has only two states. For example, if a feature/function is *enabled*, toggling changes the state to *disabled*.

topology Logical and/or physical arrangement of switches on a network.

trap An asynchronous (unsolicited) notification of an event originating on an SNMP-managed device and directed to a centralized SNMP Network Management Station.

U

unblocked port Devices communicating with an unblocked port can log in to a Connectrix switch or a similar product and communicate with devices attached to any other unblocked port if the devices are in the same zone.

Unicast Unicast routing provides one or more optimal path(s) between any of two switches that make up the fabric. (This is used to send a single copy of the data to designated destinations.)

upper layer protocol (ULP) The protocol user of FC-4 including IPI, SCSI, IP, and SBCCS. In a device driver ULP typically refers to the operations that are managed by the class level of the driver, not the port level.

URL Uniform Resource Locator, the addressing system used by the World Wide Web. It describes the location of a file or server anywhere on the Internet.

V

virtual switch A Fibre Channel switch function that allows users to subdivide a physical switch into multiple virtual switches. Each virtual switch consists of a subset of ports on the physical switch, and has all the properties of a Fibre Channel switch. Multiple virtual switches can be connected through ISL to form a virtual fabric or VSAN.

virtual storage area network (VSAN) An allocation of switch ports that can span multiple physical switches, and forms a virtual fabric. A single physical switch can sometimes host more than one VSAN.

volume A general term referring to an addressable logically contiguous storage space providing block I/O services.

VSAN Virtual Storage Area Network.

W

warning message An indication that a possible error has been detected. *See also* “[error message](#)” and “[information message](#).”

World Wide Name (WWN) A unique identifier, even on global networks. The WWN is a 64-bit number (XX:XX:XX:XX:XX:XX:XX:XX). The WWN contains an OUI which uniquely determines the equipment manufacturer. OUIs are administered by the Institute of Electronic and Electrical Engineers (IEEE). The Fibre Channel environment uses two types of WWNs; a World Wide Node Name (WWNN) and a World Wide Port Name (WWPN). Typically the WWPN is used for zoning (path provisioning function).

Z

zone An information object implemented by the distributed Nameserver (dNS) of a Fibre Channel switch. A zone contains a set of members which are permitted to discover and communicate with one another. The members can be identified by a WWPN or port ID. EMC recommends the use of WWPNs in zone management.

zone set An information object implemented by the distributed Nameserver (dNS) of a Fibre Channel switch. A Zone Set contains a set of Zones. A Zone Set is activated against a fabric, and only one Zone Set can be active in a fabric.

zonie A storage administrator who spends a large percentage of his workday zoning a Fibre Channel network and provisioning storage.

zoning Zoning allows an administrator to group several devices by function or by location. All devices connected to a connectivity product, such as a Connectrix switch, may be configured into one or more zones.

