

EMC CLARiiON
CX300, CX500, and CX700 Storage Systems

CONFIGURATION PLANNING GUIDE

P/N 300-001-273
REV A03

EMC Corporation
Corporate Headquarters:
Hopkinton, MA 01748 -9103
1-508 -435 -1000
www.EMC.com

Copyright © 2003, 2004 EMC® Corporation. All rights reserved.

Published August 2004

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.

Trademark Information

EMC², EMC, CLARiiON, Navisphere, PowerPath, and Symmetrix are registered trademarks and Access Logix, FLARE, MirrorView, Powerlink, SAN Copy, SnapView, and TimeFinder are trademarks of EMC Corporation.

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

Preface	ix
----------------------	----

Chapter 1 About CX300, CX500, and CX700 Storage Systems

Introducing CX300, CX500, and CX700 Storage Systems	1-2
Fibre Channel Background	1-4
Fibre Channel Storage Components	1-5
Server Component (Host Bus Adapter and Driver)	1-5
Interconnect Components	1-5
Storage Component (Storage System Enclosures)	1-8
Types of Storage-System Installations	1-9
About Switched Shared Storage and SANs (Storage Area Networks)	1-10
Storage Groups.....	1-11

Chapter 2 RAID Types and Trade-offs

Introducing RAID	2-2
Disk Striping.....	2-2
Mirroring.....	2-2
RAID Groups and LUNs	2-3
RAID Types.....	2-4
RAID 5 Group (Individual Access Array)	2-4
RAID 3 Group (Parallel Access Array).....	2-5
RAID 1 Mirrored Pair.....	2-7
RAID 1/0 Group (Mirrored RAID 0 Group)	2-8
RAID 0 Group (Nonredundant Array).....	2-10
Individual Disk Unit	2-10
Hot Spare	2-11

RAID Benefits and Trade-offs	2-13
Performance	2-14
Storage Flexibility	2-15
Data Availability and Disk Space Usage.....	2-16
Guidelines for RAID Groups	2-18
Sample Applications for RAID Types.....	2-20

Chapter 3 Storage-System Management

Navisphere Management Suites.....	3-2
Access Logix Software	3-2
Navisphere Manager.....	3-3
Navisphere Host Agent and CLI.....	3-3
Navisphere Analyzer	3-3
Navisphere Integrator.....	3-3
Using Navisphere Software	3-4
Navisphere Management Worksheet	3-6

Chapter 4 Data Replication and Data Mobility Software

Data Replication and Data Mobility Software Overview	4-2
What Is SnapView?	4-3
Clones Overview and Example	4-3
Snapshots Overview and Example	4-5
Clone and Snapshot Trade-offs	4-7
SnapView Worksheets	4-8
What is MirrorView?	4-11
MirrorView Example	4-11
MirrorView Features and Benefits.....	4-13
MirrorView Planning Worksheet.....	4-15
What is MirrorView/A?	4-17
MirrorView/A Example.....	4-18
MirrorView/A Features and Benefits	4-19
MirrorView/A Planning Worksheets	4-22
What Is EMC SAN Copy?	4-24
SAN Copy Example	4-26
SAN Copy Features and Benefits.....	4-27
Implementing SAN Copy Over Extended Distances	4-28
SAN Copy Planning Worksheet.....	4-30

Chapter 5	Planning File Systems and LUNs	
	Multiple Paths to LUNs.....	5-2
	Sample Shared Switched Installation.....	5-2
	Sample Unshared Direct Installation	5-5
	Planning Applications, LUNs, and Storage Groups	5-6
	Application and LUN Planning.....	5-7
	LUN and Storage Group Planning Worksheet.....	5-9
	LUN Details Worksheet	5-14
 Chapter 6	 Storage-System Hardware	
	Storage-System Installation Types	6-2
	Hardware for CX300, CX500, and CX700 Storage Systems	6-3
	Storage Hardware	6-3
	Storage Processor (SP).....	6-5
	Disks.....	6-7
	Planning Your Hardware Components	6-8
	Components for Shared Switched Storage.....	6-8
	Components for Shared-or-Clustered Direct Storage	6-8
	Components for Unshared Direct Storage	6-8
	Hardware Dimensions and Requirements	6-9
	Cabinets for Enclosures	6-10
	Data Cable and Configuration Guidelines	6-10
	Hardware Planning Worksheets	6-12
 Index		 i-1

1-1	CX700 Storage System	1-3
1-2	Nodes - Initiator and Target	1-4
1-3	Switch and Point-to-Point Connections	1-6
1-4	A Switch Zone	1-7
1-5	Typical 16-Port Switch, Back View	1-7
1-6	CX700 Storage Processor Enclosure (SPE)	1-8
1-7	Types of Storage-System Installation	1-9
1-8	Components of a SAN	1-10
1-9	Sample Shared Storage Configuration	1-12
1-10	Data Access Control with Shared Storage	1-13
2-1	Multiple LUNs in a RAID Group	2-3
2-2	RAID 5 Group	2-5
2-3	RAID 3 Group	2-6
2-4	RAID 1 Mirrored Pair	2-8
2-5	RAID 1/0 Group	2-9
2-6	How a Hot Spare Works	2-12
2-7	Disk Space Usage in the RAID Configuration	2-17
3-1	Storage Domains on the Internet	3-5
4-1	Clones Example	4-4
4-2	Snapshot Example	4-6
4-3	Sample MirrorView Configuration	4-12
4-4	Sample Remote Mirror Configuration	4-19
4-5	Sample SAN Copy Sessions Between Storage Systems	4-26
4-6	Sample SAN Copy Over IP Configuration	4-29
5-1	Sample Shared Switched Storage Configuration	5-3
5-2	Unshared Direct Storage System	5-5
6-1	Types of Storage-System Installation	6-2
6-2	CX700 Storage System	6-4
6-3	Shared Storage Systems	6-5
6-4	Disks and Their IDs	6-7

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

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

About this Manual

This planning guide introduces EMC CLARiiON® CX300, CX500, and CX700 disk-array storage systems and offers useful background information and worksheets to help you plan.

Audience

Please read this guide

- ◆ if you are considering purchase of a CX300, CX500, or CX700 disk-array storage system and want to understand its features; or
- ◆ before you plan the installation of a storage system.

You should be familiar with the host servers that will use the storage systems and with the operating systems on the servers. After reading this guide, you will be able to

- ◆ Determine the best storage-system components for your installation
- ◆ Determine your site requirements
- ◆ Configure storage systems correctly

Organization

This manual is organized as follows:

- Chapter 1 Provides background information about Fibre Channel features and explains the major types of storage.
- Chapter 2 Describes the RAID Groups and the different ways they store data.
- Chapter 3 Explains storage management software.
- Chapter 4 Describes the optional EMC storage-system based applications for replicating data or copying data within or between CLARiiON storage systems:
 - ◆ EMC SnapView™ software for replicating data with a storage system using snapshots or clones.
 - ◆ EMC MirrorView™ and EMC MirrorView / Asynchronous remote mirroring software for replicating data between storage systems.
 - ◆ EMC SAN Copy™ software for creating clones and snapshots within a storage system.
- Chapter 5 Helps you plan your storage-system file systems and LUNs.
- Chapter 6 Explains the hardware components of storage systems.

Conventions Used in This Guide

EMC uses the following conventions for notes and cautions.

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

**CAUTION**

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

Finding Current Information

The most up-to-date version of this manual is posted on the EMC Powerlink™ website. We recommend that you download the latest version before you plan the installation of a storage system.

To access EMC Powerlink, use the following link:

<http://powerlink.emc.com>

After you log in, select **Support > Document Library**.

Where to Get Help

For questions about technical support, call your local sales office or service provider.

If you have a valid EMC service contract, contact EMC Customer Service at:

United States: (800) 782-4362 (SVC-4EMC)

Canada: (800) 543-4782 (543-4SVC)

Worldwide: (508) 497-7901

Follow the voice menu prompts to open a service call and select the applicable product support.

Sales and Customer Service Contacts

For the list of EMC sales locations, please access the EMC home page at:

<http://www.EMC.com/contact/>

For additional information on the EMC products and services available to customers and partners, refer to the EMC Powerlink Web site at:

<http://powerlink.EMC.com>

Your Comments

Your suggestions will help us continue to improve the accuracy, organization, and overall quality of the user publications. Please send a message to techpub_comments@EMC.com with your opinions of this guide.

About CX300, CX500, and CX700 Storage Systems

This chapter introduces CLARiiON® CX300, CX500, and CX700 disk-array storage systems and storage area networks (SANs).

Major sections are

- ◆ Introducing CX300, CX500, and CX700 Storage Systems 1-2
- ◆ Fibre Channel Background 1-4
- ◆ Fibre Channel Storage Components..... 1-5
- ◆ Types of Storage-System Installations..... 1-9
- ◆ About Switched Shared Storage and SANs (Storage Area Networks)..... 1-10

Introducing CX300, CX500, and CX700 Storage Systems

EMC CLARiiON CX300, CX500, and CX700 disk-array storage systems provide terabytes of disk storage capacity, high transfer rates, flexible configurations, and highly available data at low cost. Their hardware RAID features are provided by two storage processors (SPs).

Table 1-1 outlines each storage system enclosure and disk capacities.

Table 1-1 Storage System Enclosures and Disk Capacities

Storage System	Enclosure Type	Number of Disks per Storage System Enclosure	Maximum Number of 2-Gigabit Disk-Array Enclosures (DAE2s) ^a and Disks
CX300	2-Gbit Disk Processor Enclosure (DPE2)	15 Fibre Channel disks	3 separate DAE2s for a total of 60 disks
CX500	2-Gbit Disk Processor Enclosure (DPE2)	15 Fibre Channel disks	7 separate DAE2s for a total of 120 disks
CX700	Storage Processor Enclosure (SPE).	An SPE does not include disks, and requires at least one 2-Gbit Disk-Array Enclosure (DAE2) with Fibre Channel disks.	16 separate DAE2s for a total of 240 disks

a.Each DAE2 enclosure has slots for 15 disks.

Two types of DAE2s are available:

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

Figure 1-1 shows a typical CX700 storage system.




Figure 1-1 CX700 Storage System

A storage-system package includes storage management software, Fibre Channel interconnect hardware, and one or more storage systems. The host bus adapter driver hardware and software are available from outside vendors.

Fibre Channel Background

Fibre Channel is a high-performance serial protocol that allows transmission of both network and I/O channel data. It is a low level protocol, independent of data types, and supports such formats as SCSI and IP.

The Fibre Channel standard supports several physical topologies, including switch fabric point-to-point and arbitrated loop (FC-AL). The Fibre Channel storage systems described in this manual use switch fabric and FC-AL topologies.

A switch fabric is a set of point-to-point connections between nodes; each connection is made through one or more Fibre Channel switches. Each node may have its own unique address, but the path between nodes is governed by a switch. The nodes are connected by optical cable.

A Fibre Channel arbitrated loop is a circuit consisting of nodes. Each node has a unique address, called a Fibre Channel arbitrated loop address. The nodes connect with optical cables. An optical cable can transmit data over great distances for connections that span entire enterprises and can support remote disaster recovery systems.

Each connected device in a switched fabric or arbitrated loop is a server adapter (initiator) or a target (storage system). The switches are not considered nodes.

Figure 1-2 shows a node and initiator.




Figure 1-2 Nodes - Initiator and Target


Fibre Channel Storage Components

A Fibre Channel storage system has three main components:

- ◆ Server component (host bus adapter driver with adapter and software)
- ◆ Interconnect components (cables based on Fibre Channel standards, and switches)
- ◆ Storage component (storage system with storage processors, power supply, cooling hardware, and disks)

Server Component (Host Bus Adapter and Driver)

The host bus adapter is a printed-circuit board that slides into an I/O slot in the server's cabinet. It transfers data between server memory and one or more disk-array storage systems over Fibre Channel — as controlled by the support software (adapter driver).



Interconnect Components

The interconnect components include the optical cables between components and any Fibre Channel switches.

The maximum length of optical cable between a CX300, CX500, and CX700 storage system and a server or switch ranges from 150 to 500 meters (165 to 550 yards), depending on the type of cable and operating speed. With extenders, connections between servers, switches, and other devices can span up to 60 kilometers (36 miles) or more. This ability to span great distances is a major advantage of optical cable.

Details on cable lengths and rules appear later, in Chapter 6.

Fibre Channel Switches

A Fibre Channel switch, which is required for switched shared storage (a storage area network, SAN), connects all the nodes cabled to it using a fabric topology. A switch adds serviceability and scalability to any installation; it allows on-line insertion and removal of any device on the fabric and maintains integrity if any connected device stops participating. A switch also provides server-to-storage-system access control and point-to-point connections.




Figure 1-3 Switch and Point-to-Point Connections

You can cascade switches (connect one switch port to another switch) for additional port connections.


Switch Zoning

Switch zoning lets an administrator define paths between connected nodes based on the node's unique World Wide Name. Each zone includes a server adapter and/or one or more SPs. EMC recommends single-initiator zoning, which limits each zone to a single HBA port (initiator). The current connection limits are outlined in Table 1-2.

Table 1-2 Connection Limits

Storage System	Maximum HBA Ports per SP	Maximum HBA Ports per Storage System
CX300	32	128
CX500	64	256
CX700	64	512

In Figure 1-4, Server 1 has access to one SP (SP A) in storage systems 1 and 2; it has no access to any other SP.




EMC1806

Figure 1-4 A Switch Zone

To illustrate switch zoning, Figure 1-4 shows just one HBA per server and one switch. Normally, such installations include multiple HBAs per server and two or more switches.

If you do not define a zone in a switch, all adapter ports connected to the switch can communicate with all SP ports connected to the switch. However, access to an SP does not necessarily provide access to the SP's storage. Access to storage is governed by the Storage Groups you create (explained later).

Fibre Channel switches are available with 8, 16, 32, or more ports. They are compact units that fit into a rackmount cabinets.



EMC1807


Figure 1-5 Typical 16-Port Switch, Back View

If your servers and storage systems will be far apart, you can place the switches closer to the servers or storage systems, as convenient.

A switch is technically a repeater, not a node, in a Fibre Channel loop. However, it is bound by the same cabling distance rules as a node.

Storage Component (Storage System Enclosures)

EMC CX300, CX500, and CX700 disk-array storage systems, with their storage processor (SPs), power supplies, and cooling hardware form the storage component of a Fibre Channel system. The controlling unit is a CX300 or CX500 disk processor enclosure (DPE2) or a CX700 storage processor enclosure (SPE). A CX700 SPE, outside of its cabinet, is shown in Figure 1-6.



EMC2434

Figure 1-6 CX700 Storage Processor Enclosure (SPE)

Each CX300 (DPE2), CX500 (DPE2), and CX700 (SPE) has two SPs. The number of ports for each SP varies between storage systems. Table 1-3 outlines the number of ports per SP for each system.

Table 1-3 Number of Storage System Ports per SP


Storage System	Front End Ports	Back End Ports
CX300	2	1
CX500	2	2
CX700	4	4

The front-end ports communicate with switches or servers and the back-end ports communicate with disks. Hardware details appear later in Chapter 6.

Types of Storage-System Installations

You can use a storage system in any of several types of installation:

- ◆ **Unshared direct** with one server is the simplest and least costly.
- ◆ **Shared-or-clustered direct**, with a limit of two servers, lets two servers share storage resources with high availability.
- ◆ **Shared switched**, with two switch fabrics, lets two or more servers share the resources of several storage systems in a storage area network (SAN). Shared switched installations are available in high-availability versions (two HBAs per server) or with one HBA per server. Shared switched storage systems can have multiple paths to each SP, providing multipath I/O for dynamic load sharing and greater throughput.



EMC1826

Figure 1-7 Types of Storage-System Installation

Storage systems for any shared installation require EMC Access Logix™ software to control server access to the storage-system LUNs.

The Shared-or-clustered direct installation can be either shared (that is, use Access Logix to control LUN access) or clustered (without Access Logix, but with operating system cluster software controlling LUN access).


About Switched Shared Storage and SANs (Storage Area Networks)

This section explains the features that let multiple servers share disk-array storage systems on a SAN (storage area network).

A SAN is one or more storage devices connected to servers through Fibre Channel switches to provide a central location for disk storage. Centralizing disk storage among multiple servers has many advantages, including

- ◆ highly available data
- ◆ flexible association between servers and storage capacity
- ◆ centralized management for fast, effective response to users' data storage needs
- ◆ easier file backup and recovery

An EMC SAN is based on shared storage; that is, the SAN requires EMC Access Logix software to provide flexible access control to storage-system LUNs. Within the SAN, a network connection to each SP in the storage system lets you configure and manage it.



EMC1810

Figure 1-8 Components of a SAN

Fibre Channel switches can control data access to storage systems through the use of switch zoning, explained earlier on page 1-6.

However, switch zoning cannot selectively control data access to LUNs in a storage system, because each SP appears as a single Fibre Channel device to the switch fabric. So switch zoning can prevent or allow communication with an SP, but not with specific disks or LUNs attached to an SP. For access control with LUNs, a different solution is required: Storage Groups.


Storage Groups

A Storage Group is one or more LUNs (logical units) within a storage system that is reserved for one or more servers and is inaccessible to other servers. Storage Groups are the central component of shared storage; storage systems that are unshared do not use Storage Groups.

When you configure shared storage, you specify servers and the Storage Group(s) each server can read from and/or write to. Access Logix software running in each storage system enforces the server-to-Storage Group permissions.

More than one server can access a Storage Group if all the servers run cluster software. The cluster software enforces orderly access to the shared Storage Group LUNs.

Figure 1-9 shows a simple shared storage configuration consisting of one storage system with two Storage Groups. One Storage Group serves a cluster of two servers running the same operating system, and the other Storage Group serves a UNIX® database server. Each server is configured with two independent paths to its data, including separate host bus adapters, switches, and SPs, so there is no single point of failure for access to its data.



EMC1811n

Figure 1-9 Sample Shared Storage Configuration

Access Control with Shared Storage

Access control permits or restricts a server's access to shared storage. Configuration access, the ability to configure storage systems, is governed by the storage system software security mechanism: a set of privileges based on username and password that's stored, encrypted, on one or more storage systems in a network.

Data access, the ability to read and write information to storage-system LUNs, is provided by Storage Groups. During storage-system configuration, using a management utility, the system administrator associates a server with one or more LUNs. The associated LUNs compose a Storage Group.

Each server sees its Storage Group as if it were an entire storage system, and never sees the other LUNs on the storage system. Therefore, it cannot access or modify data on LUNs that are not part of its Storage Group. However, you can define a Storage Group as sharable (accessible by more than one server), if, as shown above in Figure 1-9, the servers run cluster software.

Figure 1-10 shows access control through Storage Groups. Each server has exclusive read and write access to its designated Storage Group.




Figure 1-10 Data Access Control with Shared Storage

What Next?

For information about RAID types and RAID trade-offs, continue to the next chapter.

For information on the Navisphere Manager Suite software, go to Chapter 3

For information on the optional data replication software (SnapView™, MirrorView™, MirrorView / Asynchronous) or data mobility software (SAN Copy™), go to Chapter 4.

To plan LUNs and file systems, go to Chapter 5.

For details on the storage-system hardware, go to Chapter 6.

This chapter explains the RAID types you can choose for your storage-system LUNs. If you already know about RAID types and know which ones you want, you can skip this background information and go to the LUN planning chapter (Chapter 5).

Major sections are

- ◆ Introducing RAID2-2
- ◆ RAID Types2-4
- ◆ RAID Benefits and Trade-offs.....2-13
- ◆ Guidelines for RAID Groups.....2-18
- ◆ Sample Applications for RAID Types2-20

Introducing RAID

The storage system uses RAID (redundant array of independent disks) technology. RAID technology groups separate disks into one logical unit (LUN) to improve reliability and/or performance.

The storage system supports five RAID levels and two other disk configurations, the individual unit and the hot spare (global spare). You group the disks into one RAID Group by *binding* them using a storage-system management utility.

Four of the RAID levels use *disk striping* and two use *mirroring*.

Disk Striping

Using disk stripes, the storage-system hardware can read from and write to multiple disks simultaneously and independently. By allowing several read/write heads to work on the same task at once, disk striping can enhance performance. The amount of information read from or written to each disk makes up the stripe element size. The stripe size is the stripe element size multiplied by the number of disks in a group. For example, assume a stripe element size of 128 sectors (the default). If the group has five disks, you would multiply five by the stripe element size of 128 to yield a stripe size of 640 sectors.

The storage system uses disk striping with most RAID types.

Mirroring

Mirroring maintains a copy of a logical disk *image* that provides continuous access if the original image becomes inaccessible. The system and user applications continue running on the good image without interruption. There are two kinds of mirroring: hardware mirroring, in which the SP synchronizes the disk images; and software mirroring, in which the operating system synchronizes the images. Software mirroring consumes server resources, since the operating system must mirror the images, and has no offsetting advantages; we mention it here only for historical completeness.

With a storage system, you can create a hardware mirror by binding disks as a RAID 1 mirrored pair or a RAID 1/0 Group (a mirrored RAID 0 Group); the hardware will then mirror the disks automatically.

With a LUN of any RAID type, a storage system can maintain a remote copy using the optional MirrorView or MirrorView / Asynchronous data replication software. This software, primarily useful for disaster recovery, is explained in Chapter 4.

RAID Groups and LUNs

Some RAID types let you create multiple LUNs on one RAID Group. You can then allot each LUN to a different user, server, or application. For example, a five-disk RAID 5 Group that uses 73-Gbyte disks offers 292 Gbytes of space. You could bind three LUNs, with 60, 100, and 132 Gbytes of storage capacity, for temporary, mail, and customer files.

One disadvantage of multiple LUNs on a RAID Group is that I/O to each LUN may affect I/O to the others in the group; that is, if traffic to one LUN is very heavy, I/O performance with other LUNs may be degraded. The main advantage of multiple LUNs per RAID Group is the ability to divide the enormous amount of disk space provided by RAID Groups on newer, high-capacity disks.

Figure 2-1 shows three LUNs in one RAID Group.




Figure 2-1 Multiple LUNs in a RAID Group

RAID Types

You can choose from the following RAID types: RAID 5, RAID 3, RAID 1, RAID 0, RAID 1/0, individual disk unit, and hot spare.

You can choose an additional type of redundant disk — a remote mirror — for any RAID type except a hot spare.

RAID 5 Group (Individual Access Array)

A RAID 5 Group usually consists of five disks (but can have three to sixteen). A RAID 5 Group uses disk striping. With a RAID 5 group, you can create up to 128 RAID 5 LUNs to apportion disk space to different users, servers, and applications.

The storage system writes parity information that lets the Group continue operating if a disk fails. When you replace the failed disk, the SP rebuilds the group using the information stored on the working disks. Performance is degraded while the SP rebuilds the group. However, the storage system continues to function and gives users access to all data, including data stored on the failed disk.

Figure 2-2 shows user and parity data with the default stripe element size of 128 sectors (65,536 bytes) in a five-disk RAID 5 group. The stripe size comprises all stripe elements. Notice that the disk block addresses in the stripe proceed sequentially from the first disk to the second, third, and fourth, then back to the first, and so on.




Figure 2-2 RAID 5 Group

RAID 5 Groups offer excellent read performance and good write performance. Write performance benefits greatly from storage-system caching.

RAID 3 Group (Parallel Access Array)

A RAID 3 Group consists of five or nine disks. The hardware always reads from or writes to all the disks. A RAID 3 Group uses disk striping. Limit concurrent access to RAID 3 LUNs for the best bandwidth performance. For example, A RAID 3 Group may have multiple LUNs, but the highest bandwidth is achieved with 1 to 4 threads of concurrent, large I/O.

The storage system writes parity information that lets the Group continue operating if a disk fails. When you replace the failed disk, the SP rebuilds the group using the information stored on the working disks. Performance is degraded while the SP rebuilds the group. However, the storage system continues to function and gives users access to all data, including data stored on the failed disk.

Figure 2-3 shows user and parity data with a data block size of 2 Kbytes in a RAID 3 Group. Notice that the byte addresses proceed from the first disk to the second, third, and fourth, then the first, and so on.




Figure 2-3 RAID 3 Group

RAID 3 differs from RAID 5 in one major way. With a RAID 3 Group, the parity information is stored on one disk; with a RAID 5 Group, it is stored on all disks. RAID 3 can perform sequential I/O better than RAID 5, but does not handle random access as well. FLARE™ Operating Environment 2.07 improved RAID 3 processing in several ways:

- ◆ Use of the write cache is now equivalent to RAID 5
- ◆ I/O sizes to disk have been increased
- ◆ Maximum host transfer size has been increased.

RAID 3 is best thought of as a specialized RAID 5, where the expectation is either large or sequential I/O. However, with write cache active for RAID 3, RAID 3 is more flexible than in previous releases, and can handle some level of concurrency. A RAID 3 Group works well for applications that use I/Os of blocks 64 Kbytes and

larger. By using both the read and write cache, a RAID 3 group can handle several concurrent streams of access.

RAID 3 Groups do not require any special buffer area. No fixed memory is required to use write cache with RAID 3. Simply allocate write cache as you would for RAID 5, and ensure that caching is turned on for the LUNs in the RAID 3 Groups. Access to RAID 3 LUNs is compatible with concurrent access to LUNs of other RAID types on the storage system. In fact, with the advent of backup to disk with the CLARiiON Data Library and ATA drives, this should be considered the usual case.

RAID 1 Mirrored Pair

A RAID 1 Group consists of two disks that are mirrored automatically by the storage-system hardware. With a RAID 1 Group, you can create multiple RAID 1 LUNs to apportion disk space to different users, servers, and applications.

RAID 1 hardware mirroring within the storage system is not the same as software mirroring, remote mirroring, or hardware mirroring for other kinds of disks. Functionally, the difference is that you cannot manually stop mirroring on a RAID 1 mirrored pair, and then access one of the images independently. If you want to use one of the disks in such a mirror separately, you must unbind the mirror (losing all data on it), rebind the disk as the type you want, and software format the newly bound LUN.

With a storage system, RAID 1 hardware mirroring has the following advantages:

- ◆ automatic operation (you do not have to issue commands to initiate it)
- ◆ physical duplication of images
- ◆ a rebuild period that you can select during which the SP recreates the second image after a failure

With a RAID 1 mirrored pair, the storage system writes the same data to both disks, as follows.




Figure 2-4 RAID 1 Mirrored Pair

RAID 1/0 Group (Mirrored RAID 0 Group)

A RAID 1/0 Group consists of four, six, eight, ten, twelve, fourteen, or sixteen disks. These disks make up two mirror images, with each image including two to eight disks. The hardware automatically mirrors the disks. A RAID 1/0 Group uses disk striping. It combines the speed advantage of RAID 0 with the redundancy advantage of mirroring. With a RAID 1/0 Group, you can create up to 128 RAID 1/0 LUNs to apportion disk space to different users, servers, and applications.

Figure 2-5 shows the distribution of user data with the default stripe element size of 128 sectors (65,536 bytes) in a six-disk RAID 1/0 Group. Notice that the disk block addresses in the stripe proceed sequentially from the first mirrored disks (first and fourth disks) to the second mirrored disks (second and fifth disks), to the third mirrored disks (third and sixth disks), and then from the first mirrored disks, and so on.




Figure 2-5 RAID 1/0 Group

A RAID 1/0 Group can survive the failure of multiple disks, providing that one disk in each image pair survives.

RAID 0 Group (Nonredundant Array)



CAUTION

A RAID 0 Group provides no protection for your data. EMC does not recommend using a RAID 0 Group unless you have some way of protecting your data, such as software mirroring.

A RAID 0 Group consists of three to a maximum of sixteen disks. A RAID 0 Group uses disk striping, in which the hardware writes to or reads from multiple disks simultaneously. You can create up to 128 LUNs in any RAID 0 Group.

Unlike the other RAID levels, with RAID 0 the hardware does not maintain parity information on any disk; this type of group has no inherent data redundancy. As a result, if any failure (including an unrecoverable read error) occurs on a disk in the LUN, the information on the LUN is lost.

RAID 0 offers enhanced performance through simultaneous I/O to different disks. A desirable alternative to RAID 0 is RAID 1/0, which does protect your data.

Individual Disk Unit

An individual disk unit is a disk bound to be independent of any other disk in the cabinet. An individual unit has no inherent high availability, but you can make it highly available by using software mirroring with another individual unit. You can create one LUN per individual disk unit. If you want to apportion the disk space, you can do so using partitions, file systems, or user directories.

Hot Spare

A hot spare is a dedicated replacement disk on which users cannot store information. A hot spare is global: if any disk in a RAID 5 Group, RAID 3 Group, RAID 1 mirrored pair, or RAID 1/0 Group fails, the SP automatically rebuilds the failed disk's structure on the hot spare. When the SP finishes rebuilding, the disk group functions as usual, using the hot spare instead of the failed disk. When you replace the failed disk, the SP copies the data from the former hot spare onto the replacement disk.


When the copy is done, the disk group consists of disks in the original slots, and the SP automatically frees the hot spare to serve as a hot spare again. A hot spare is most useful when you need the highest data availability. It eliminates the time and effort needed for someone to notice that a disk has failed, find a suitable replacement disk, and insert the disk.

When you plan to use a hot spare, make sure the disk has the capacity to serve in any RAID Group in the storage-system chassis. A RAID Group cannot use a hot spare that is smaller than a failed disk in the group.

You can have one or more hot spares per storage-system chassis. You can make any disk in the chassis a hot spare, except a disk that stores the FLARE™ operating environment (formerly called Base or Core Software) or the write cache vault; that is, a hot spare can be any disk except disk IDs 000-004.

If you use hot spares of different sizes, the storage system will automatically use the hot spare of the proper size in place of a failed disk.

An example of hot spare usage for a CX700 storage system follows.



1. RAID 5 Groups consist of disk modules 0-4 and 5-9; mirrored pairs are modules 10-11 and 12-13, disk module 14 is a hot spare.
2. Disk module 3 fails
3. RAID 5 Group becomes modules 0, 1, 2, 14, and 4; now no hot spare is available.
4. System operator replaces failed module 3 with a functional module.
5. Once again, RAID 5 Group consists of modules 0-4 and the hot spare is 14.

EMC2432

Figure 2-6 How a Hot Spare Works

RAID Benefits and Trade-offs

This section reviews RAID types and explains their benefits and trade-offs. You can create seven types of LUN:

- ◆ RAID 5 Group (individual access array)
- ◆ RAID 3 Group (parallel access array)
- ◆ RAID 1 mirrored pair
- ◆ RAID 1/0 Group (mirrored RAID 0 Group); a RAID 0 Group mirrored by the storage-system hardware
- ◆ RAID 0 Group (nonredundant individual access array); no inherent high-availability features
- ◆ Individual unit; no inherent high-availability features
- ◆ Hot spare; serves only as an automatic replacement for any disk in a RAID type other than 0; does not store data during normal system operations

Plan the disk unit configurations carefully. After a disk has been bound into a LUN, you cannot change the RAID type of that LUN without unbinding it, and this means losing all data on it.

Table 2-1 compares the read and write performance and relative cost per gigabyte (Gbyte) of the RAID types. Figures shown are theoretical maximums.

Table 2-1 Performance, Availability, and Cost of RAID Types (Individual Unit = 1.0)

Disk Configuration	Relative Read Performance Without Cache	Relative Write Performance Without Cache	Relative Cost Per Gbyte
RAID 5 Group with five disks	Up to 5x with five disks (for small I/O requests that are 2 to 8 Kbytes)	Up to 1.25x with five disks (for small I/O requests that are 2 to 8 Kbytes)	1.25
RAID 3 Group with five disks	Up to 4x (for large I/O requests)	Up to 4x (for large I/O requests)	1.25
RAID 1 mirrored pair	Up to 2x	Up to 1x	2
RAID 1/0 Group with 10 disks	Up to 10x	Up to 5x	2
Individual unit	1x	1x	1
Notes: These performance numbers are not based on storage-system caching. With caching, the performance numbers for RAID 5 writes improve significantly. Performance multipliers vary with the load on the server and storage system.			

Performance

RAID 5, with individual access, provides high read throughput by allowing simultaneous reads from each disk in the group. RAID 5 write performance is excellent when the storage system uses write caching.

RAID 3, with parallel access, provides high throughput for sequential, large block-size requests (blocks of more than 64 Kbytes). With RAID 3, the system accesses all disks in each request but need not read data and parity before writing – advantageous for large requests but not for small ones. RAID 3 uses SP memory without caching, which means you do not need the second SP and SPS that caching requires.

Generally, the performance of a RAID 3 Group increases as the size of the I/O request increases. Read performance increases rapidly with read requests up to 1 Mbyte. Write performance increases greatly for sequential write requests that are greater than 256 Kbytes. For applications issuing very large I/O requests, a RAID 3 LUN provides significantly better write performance than a RAID 5 LUN.

We do not recommend using RAID 3 in the same storage-system chassis with RAID 5 or RAID 1/0.

A RAID 1 mirrored pair has its disks locked in synchronization, but the SP can read data from the disk whose read/write heads are closer to it. Therefore, RAID 1 read performance can be twice that of an individual disk while write performance remains the same as that of an individual disk.

A RAID 0 Group (nonredundant individual access array) or RAID 1/0 Group (mirrored RAID 0 Group) can have as many I/O operations occurring simultaneously as there are disks in the group. Since RAID 1/0 locks pairs of RAID 0 disks the same way as RAID 1 does, the performance of RAID 1/0 equals the number of disk pairs times the RAID 1 performance number. If you want high throughput for a specific LUN, use a RAID 1/0 or RAID 0 Group. A RAID 1/0 Group requires at least four disks; a RAID 0 Group, at least three disks.

An individual unit has only one I/O operation per read or write operation.

RAID types 5, 1, 1/0, and 0 allow multiple LUNs per RAID Group. If you create multiple LUNs on a RAID Group, the LUNs share the RAID Group disks, and the I/O demands of each LUN affect the I/O service time to the other LUNs. For best performance, you may want to use one LUN per RAID Group.

Storage Flexibility

Certain RAID Group types — RAID 5, RAID 1, RAID 1/0, and RAID 0 — let you create up to 128 LUNs in each group. This adds flexibility, particularly with large disks, since it lets you apportion LUNs of various sizes to different servers, applications, and users. Conversely, with RAID 3, there can be only one LUN per RAID Group, and the group must include five or nine disks — a sizable block of storage to devote to one server, application, or user. However, the nature of RAID 3 makes it ideal for that single-threaded type of application.

Data Availability and Disk Space Usage


If data availability is critical and you cannot afford to wait hours to replace a disk, rebind it, make it accessible to the operating system, and load its information from backup, then use a redundant RAID Group: RAID 5, RAID 3, RAID 1 mirrored pair, or RAID 1/0. If data availability is not critical, or disk space usage is critical, bind an individual unit.

A RAID 1 mirrored pair or RAID 1/0 Group provides very high data availability. They are more expensive than RAID 5 or RAID 3 Groups, since only 50 percent of the total disk capacity is available for user data, as shown in Figure 2-7 on page 2-17.

A RAID 5 or RAID 3 Group provides high data availability, but requires more disks than a mirrored pair. In a RAID 5 or RAID 3 Group, the disk space available for user data is the total capacity of one less than the number of disks in the RAID Group. For example, in a five-disk RAID 5 Group, the capacity of four disks is available for user data, which is 80% of the group's total disk capacity. So RAID 5 and RAID 3 Groups use disk space much more efficiently than a mirrored pair. A RAID 5 or RAID 3 Group is usually more suitable than a RAID 1 mirrored pair for applications where high data availability, good performance, and efficient disk space usage are all of relatively equal importance.

For a LUN in any RAID Group, you can provide for disaster recovery by establishing a remote mirror at a distant site.

Figure 2-7 illustrates RAID type disk space usage.



EMC1820

Figure 2-7 Disk Space Usage in the RAID Configuration

A RAID 0 Group (nonredundant individual access array) provides all its disk space for user files, but does not provide any high availability features. For high availability, you can use a RAID 1/0 Group instead.

A RAID 1/0 Group provides the best combination of performance and availability, at the highest cost per Gbyte of disk space.

An individual unit, like a RAID 0 Group, provides no high-availability features. All its disk space is available for user data, as shown in Figure 2-7 above.

Guidelines for RAID Groups

To decide when to use a RAID 5 Group, RAID 3 Group, mirror (that is, a RAID 1 mirrored pair or RAID 1/0 Group, a RAID 0 Group, individual disk unit, or hot spare), you need to weigh these factors:

- ◆ Importance of data availability
- ◆ Importance of performance
- ◆ Amount of data stored
- ◆ Cost of disk space

The following guidelines will help you decide on RAID types.

Use a RAID 5 Group (individual access array) for applications where

- ◆ Data availability is very important.
- ◆ Large volumes of data will be stored.
- ◆ Multitask applications use I/O transfers of different sizes.
- ◆ Excellent read and good write performance is needed (write performance is very good with write caching).
- ◆ You want the flexibility of multiple LUNs per RAID Group.

Use a RAID 3 Group (parallel access array) for applications where

- ◆ Data availability is very important.
- ◆ Large volumes of data will be stored.
- ◆ A single-task application uses large I/O transfers (more than 64 Kbytes). The operating system must allow transfers aligned to start at disk addresses that are multiples of 2 Kbytes from the start of the LUN.

Use a RAID 1 mirrored pair for applications where

- ◆ Data availability is very important.
- ◆ Speed of write access is important and write activity is heavy.

Use a RAID 1/0 Group (mirrored nonredundant array) for applications where

- ◆ Data availability is critically important.
- ◆ Overall performance is very important.

Use a RAID 0 Group (nonredundant individual access array) for applications where

- ◆ High availability is not important.
- ◆ You can afford to lose access to all data stored on a LUN if a single disk fails.
- ◆ Overall performance is very important.

Use an individual unit for applications where

- ◆ High availability is not important.
- ◆ Speed of write access is somewhat important.

Use a hot spare where

- ◆ In any RAID 5, RAID 3, RAID 1/0 or RAID 1 Group, high availability is so important that you want to regain data redundancy quickly without human intervention if any disk in the group fails.
- ◆ Minimizing the degraded performance caused by disk failure in a RAID 5 or RAID 3 Group is important.

Sample Applications for RAID Types

This section describes some types of applications in which you would want to use a RAID 5 Group, RAID 3 Group, RAID 1 mirrored pair, RAID 0 Group (nonredundant array), RAID 1/0 Group, or individual unit.

RAID 5 Group (individual access array) — Useful as a database repository or a database server that uses a normal or low percentage of write operations (writes are 33 percent or less of all I/O operations). Use a RAID 5 Group where multitasking applications perform I/O transfers of different sizes. Write caching can significantly enhance the write performance of a RAID 5 Group.

For example, a RAID 5 Group is suitable for multitasking applications that require a large history database with a high read rate, such as a database of legal cases, medical records, or census information. A RAID 5 Group also works well with transaction processing applications, such as an airline reservations system, where users typically read the information about several available flights before making a reservation, which requires a write operation. You could also use a RAID 5 Group in a retail environment, such as a supermarket, to hold the price information accessed by the point-of-sale terminals. Even though the price information may be updated daily, requiring many write operations, it is read many more times during the day.

RAID 3 Group — A RAID 3 Group (parallel access array) works well with a single-task application that uses large I/O transfers (more than 64 Kbytes), aligned to start at a disk address that is a multiple of 2 Kbytes from the beginning of the logical disk. RAID 3 Groups can use SP memory to great advantage without the second SP and battery backup unit required for storage-system caching.

You might use a RAID 3 Group for a single-task application that does large I/O transfers, like a weather tracking system, geologic charting application, medical imaging system, or video storage application.

RAID 1 mirrored pair — A RAID 1 mirrored pair is useful for logging or record-keeping applications because it requires fewer disks than a RAID 0 Group (nonredundant array) and provides high availability and fast write access. Or you could use it to store daily updates to a database that resides on a RAID 5 Group, and then, during off-peak hours, copy the updates to the database on the RAID 5 Group.

RAID 0 Group (nonredundant individual access array) — Use a RAID 0 Group where the best overall performance is important. A RAID 0 Group (like a RAID 5 Group) requires a minimum of three disks. A RAID 0 Group is useful for applications using short-term data to which you need quick access.

RAID 1/0 Group (mirrored RAID 0 Group) — A RAID 1/0 Group provides the best balance of performance and availability. You can use it very effectively for any of the RAID 5 applications. A RAID 1/0 Group requires a minimum of four disks.

Individual unit — An individual unit is useful for print spooling, user file exchange areas, or other such applications, where high availability is not important or where the information stored is easily restorable from backup.

The performance of an individual unit is slightly less than a standard disk not in an storage system. The slight degradation results from SP overhead.

Hot spare — A hot spare provides no data storage but enhances the availability of each RAID 5, RAID 3, RAID 1, and RAID 1/0 Group in a storage system. Use a hot spare where you must regain high availability quickly without human intervention if any disk in such a RAID Group fails. A hot spare also minimizes the period of degraded performance after a RAID 5 or RAID 3 disk fails.

What Next?

For information on the Navisphere Manager Suite software, go to Chapter 3.

For information on the optional data replication software (SnapView™, MirrorView™, MirrorView / Asynchronous) or data mobility software (SAN Copy™), go to Chapter 4.

To plan LUNs and file systems, go to Chapter 5.

For details on the storage-system hardware, go to Chapter 6.

Storage-System Management

This chapter explains the applications you can use to manage storage systems from servers.

Major sections are

- ◆ Navisphere Management Suites3-2
- ◆ Using Navisphere Software.....3-4

Navisphere Management Suites

You can choose one of three different Navisphere management suites: Navisphere Workgroup; Navisphere Departmental, and Navisphere Enterprise. Each suite provides the following software:

- ◆ Access Logix
- ◆ Navisphere Manager
- ◆ Navisphere Host Agents and host-based command line interpreter (CLI) for any operating system
- ◆ Navisphere Analyzer, provided with Navisphere Enterprise and an optional product you can purchase for Navisphere Workgroup and Navisphere Departmental.

The maximum number of servers that you can connect to a server varies with the management suite as follows:

Management Suite	Maximum Number of Servers Connected to a Storage System
Workgroup	2
Departmental	7
Enterprise	Maximum number defined for the specific storage-system type

Access Logix Software

Access Logix lets you restrict server access to LUNs through the use of Storage Groups. A server can access only the LUNs in the Storage Group that you connect to the server. When you create a Storage Group, you select the LUNs it contains. If multiple servers are connected to a storage system, you should connect them to different Storage Groups, unless they are part of a cluster.

Navisphere Manager

Navisphere Manager (called Manager) lets you manage multiple storage systems on multiple servers simultaneously. It includes an Event Monitor that checks storage systems for fault conditions and can notify you and/or customer service if any fault condition occurs.

With Manager you can create RAID groups, LUNs, and Storage Groups, manipulate caches, examine storage-system status and logged events, transfer control from one SP to another, and examine events recorded in storage-system event logs.

Manager features a user interface (UI) with extensive online help. UIs and online help for each of the data replication, data mobility, and Analyzer software integrate seamlessly into the Manager UI and online help.

Navisphere Host Agent and CLI

Navisphere Host Agents are provided for different operating systems. Each server connected to a storage system must run the Host Agent for its operating system. The Host Agent registers the host HBAs with the storage system and provides the storage system with information on how the server maps devices to LUNs (for example, the Windows drive letter for a LUN).

The CLI provides commands that implement Navisphere Manager UI functions in addition to commands that implement the functions of the UIs for the optional data replication and data mobility software. A major benefit offered by the CLI is the ability to write command scripts to automate management operations.

Navisphere Analyzer

Navisphere Analyzer lets you measure, compare, and chart the performance of SPs, LUNs, and disks.

Navisphere Integrator

Navisphere Integrator provides an interface between Navisphere products and HP OpenView®, CA Unicenter®, and Tivoli® network management software.

Using Navisphere Software


With Manager you can assign storage systems on an intranet or internet to a storage domain. For any installation, you can create one or more domains, provided that each storage system is in only one domain. Each storage domain must have at least one member with installed Manager.

All CX-Series storage systems now ship from the factory with Manager installed on them.

Each storage system in the domain is accessible from any other in the domain. Using an internet browser, you point at a storage system that has Manager installed. The security software then prompts you to log in. After logging in, depending on the privileges of your account, you can monitor, manage, and/or define user accounts for any storage system in the domain. You cannot view storage systems outside the domain from within the domain.

You can run the internet browser on any supported PC or laptop with a network controller and browser software. At least one storage system — ideally at least two for higher availability — in a domain must have Manager installed.

Figure 3-1 shows an internet that connects nine storage systems. It shows two domains, a U.S. Division domain with five storage systems (four systems on SANs) and a European Division with four storage systems. The 13 servers that use the storage may be connected to the same or a different network, but the intranet shown is the one used to manage the storage systems.



EMC2276

Figure 3-1 Storage Domains on the Internet

Navisphere Management Worksheet

The following worksheet will help you plan your Navisphere storage management configuration. A sample worksheet follows the blank worksheet.

Storage Management Worksheet for Navisphere

Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐
Storage System Name _____	Operating System _____	Domain Name _____
Navisphere Workgroup ☐	Navisphere Departmental ☐	Navisphere Enterprise ☐ Analyzer ☐

Sample Storage Management Worksheet for Navisphere

Write the hostnames, operating systems, and connected storage system numbers of all servers that will use the storage.

Server: cpc1033 Op sys: Windows 2000 Stor sys number/name 1
 Server: cpc1034 Op sys: Windows 2000 Stor sys number/name 1
 Server: cpc1035 Op sys: Windows 2000 Stor sys number/name 1
 Server: cpc1036 Op sys: Windows 2000 Stor sys number/name 2 3
 Server: cpc1040 Op sys: Solaris Stor sys number/name 4
 Server: cpc1041 Op sys: Solaris Stor sys number/name 4
 Server: cpc1042 Op sys: Solaris Stor sys number/name 5
 Server: cpc1043 Op sys: Linux Stor sys number/name 5

Write the domain number/name, storage system number/name (number corresponding to list above), and mark the appropriate product box. One domain can have many storage systems and needs at least one copy of a Navisphere product.

Domain number/name: 1 Stor sys number/name 1 Navisphere products: Manager ☒ Analyzer ☐
 Domain number/name: 1 Stor sys number/name 2 Navisphere products: Manager ☐ Analyzer ☒
 Domain number/name: 1 Stor sys number/name 3 Navisphere products: Manager ☐ Analyzer ☐
 Domain number/name: 2 Stor sys number/name 4 Navisphere products: Manager ☐ Analyzer ☐
 Domain number/name: 2 Stor sys number/name 5 Navisphere products: Manager ☒ Analyzer ☒
 Domain number/name: _____ Stor sys number/name _____ Navisphere products: Manager ☐ Analyzer ☐
 Domain number/name: _____ Stor sys number/name _____ Navisphere products: Manager ☐ Analyzer ☐
 Domain number/name: _____ Stor sys number/name _____ Navisphere products: Manager ☐ Analyzer ☐

What Next?

For information on the optional data replication software (SnapView™, MirrorView™, MirrorView / Asynchronous) or data mobility software (SAN Copy™), go to Chapter 4.

To plan LUNs and file systems, go to Chapter 5.

For details on the storage-system hardware, go to Chapter 6.

Data Replication and Data Mobility Software

This chapter introduces EMC CLARiiON data replication software (SnapView™, MirrorView™, and MirrorView/Asynchronous software) and the EMC data mobility software (SAN Copy™ software) for the CX500 and CX700 storage systems. These software applications allow you to create copies of logical units (LUNs).

This chapter refers to the EMC MirrorView/Asynchronous software as MirrorView/A.

Major sections are

Data Replication and Mobility Software Summary

- ◆ Data Replication and Data Mobility Software Overview4-2

SnapView

- ◆ What Is SnapView?4-3
- ◆ Clones Overview and Example4-3
- ◆ Snapshots Overview and Example.....4-5
- ◆ Clone and Snapshot Trade-offs4-7
- ◆ SnapView Worksheets4-8

MirrorView

- ◆ What is MirrorView? 4-11
- ◆ MirrorView Example 4-11
- ◆ MirrorView Features and Benefits.....4-13
- ◆ MirrorView Planning Worksheet4-15

MirrorView/Asynchronous

- ◆ What is MirrorView / A? 4-17
- ◆ MirrorView / A Example 4-18
- ◆ MirrorView / A Features and Benefits 4-19
- ◆ MirrorView / A Planning Worksheets 4-22

SAN Copy

- ◆ What Is EMC SAN Copy? 4-24
- ◆ Implementing SAN Copy Over Extended Distances 4-28
- ◆ MirrorView / A Features and Benefits 4-19
- ◆ MirrorView / A Planning Worksheets 4-22

Data Replication and Data Mobility Software Overview

Table 4-1 summarizes the common use, supported storage systems, and requirements for the data replication and data mobility software explained in this chapter.

Table 4-1 Data Replication and Data software Summary

	Data Replication			Data Mobility
	SnapView	MirrorView	MirrorView/A	SAN Copy
Common Usage	- Local backup and recovery - Application testing	- Disaster recovery - Remote backup - limited distance	- Disaster recovery - Remote backup - unlimited distance	- Data distribution - Data migration - Remote backup - unlimited distance
Supported Storage Systems	CLARiiON only	CLARiiON only	CLARiiON only	CLARiiON, Symmetrix and third party storage systems.
Requirements	Software must be installed on the CLARiiON storage system.	Software must be installed on every CLARiiON storage system with LUNs in a remote mirror.	Software must be installed on every CLARiiON storage system with LUNs in a remote mirror.	Software must be installed on at least one CLARiiON storage system participating in the SAN Copy session.^a

- a. For a full copy, SAN Copy must be installed on either the storage system with source logical or the storage system with the target logical unit. For an incremental copy, SAN Copy must be installed on the storage system with the source logical unit.

What Is SnapView?

SnapView is storage-system-based software that allows you to create a copy of a LUN by using either clones or snapshots.

A clone (BCV) is an actual copy of a LUN and takes time to create, depending on the size of the source LUN. A snapshot is a virtual point-in-time copy of a LUN and takes only seconds to create.

SnapView has the following important benefits:

- ◆ It allows full access to production data with modest impact on performance and without the risk of damaging the original data.
- ◆ For decision support or revision testing, it provides a coherent, readable and writable copy of real production data.
- ◆ For backup, it practically eliminates the time that production data spends offline or in hot backup mode, and it offloads the backup overhead from the production host to another host.

Depending on your application needs, you can create clones, snapshots, or snapshots of clones. For a detailed overview of clones or snapshots, refer to the *Clones Overview and Example* (see page 4-3) or to the *Snapshots Overview and Example* (see page 4-5). For a comparison of using clones, snapshots, and snapshots of clones, refer to Table 4-2, *A Comparison of Clones and Snapshots*, on page 4-7.

Clones Overview and Example

A clone is a complete copy of a source LUN. You specify a source LUN when you create a Clone Group. The copy of the source LUN begins when you add a clone LUN to the Clone Group. The software assigns each clone a clone ID. This ID remains with the clone until you remove the clone from its group.

While the clone is part of the Clone Group and unfractured (not accessible to a second host), any host write requests made to the source LUN are copied to the clone. Once the clone contains the required data, you can fracture the clone. Fracturing the clone breaks it from its source LUN and, once activated, makes it available to a second host.

Clone private LUNs record information that identifies data chunks on the source LUN and clone LUN that have been modified after you fractured the clone. A modified data chunk is a chunk of data that a host changes by writing to the clone or source LUN. A log in the clone private LUN records this information, but no actual data is written to the clone private LUN. This log reduces the time it takes to synchronize or reverse synchronize a clone and its source LUN since the software copies only modified (changed) chunks.

Host I/O requests made to the source LUN are copied to the fractured clone only when you manually perform a synchronization, which unfractures the clone and updates the contents on the clone with its source.

Figure 4-1 shows an example of how a fractured clone works. Note, as the production host writes to the source LUN, and the second host writes to the clone, the clone private LUN records information that identifies these writes.




Figure 4-1 Clones Example

Snapshots Overview and Example

A snapshot is a virtual LUN that allows a second host to view a point-in-time copy of a source LUN. You determine the point in time when you start a SnapView session. The session keeps track of how the source LUN looks at a particular point in time.

During a session, the production host is still able to write to the source LUN and modify data. When this happens, the software stores a copy of the original point-in-time data on a reserved LUN in the SP's LUN pool. This operation is referred to as copy-on-first-write because it occurs only when a data chunk is first modified on the source LUN.

As the session continues and additional I/O modifies data chunks on the source LUN, the amount of data stored in the reserved LUN pool grows. If needed, you can increase the size of the reserved LUN pool by adding more LUNs to the LUN pool.

Important An adequate number of reserved LUNs is essential since SnapView terminates sessions if the reserved LUN runs out of space and no additional LUNs are in the SP's LUN pool.

From a second host, you can view a session's point-in-time data by activating (mapping) a snapshot to the session. You can activate only one snapshot at a time to a session. If another point-in-time view is desired, you can deactivate (unmap) a snapshot from a session and activate it to another session of the same source LUN.

A snapshot is a composite of data chunks residing on the source LUN and on the reserved LUN in the SP's LUN pool. The data chunks on the source LUN have not been modified since you started the session. The data chunks in the reserved LUN pool are copies of the original source LUN data chunks that have been modified since you started the session. Though a snapshot appears as a conventional LUN to other hosts, its data does not reside on a disk like a conventional LUN.

Since a snapshot does not actually copy data (this is done by the session), but only views the composite of the source LUN data and the reserved LUN data, you can create or destroy a snapshot in seconds regardless of the size of the source LUN.

SnapView also allows you to instantly restore a session's point-in-time data back to the source LUN, if the source LUN were to become corrupt or if a session's point-in-time data is desired as the source. You can do this by using SnapView's rollback feature.

Figure 4-2 shows an example of how snapshots work. The snapshot is a composite of the source LUN and the reserved LUN in the SP's LUN pool.




Figure 4-2 Snapshot Example

Clone and Snapshot Trade-offs

The following table describes the benefits and trade-offs of using clones, snapshots, and snapshots of clones.

Table 4-2 A Comparison of Clones and Snapshots

	Clones	Snapshots	Snapshots of Clones
Benefits	- Provides enhanced protection against critical data loss because it is an actual LUN. - Provides immediacy in replacing the contents of the source LUN with the contents of the clone LUN, should the source become corrupted. - Provides immediacy in redirecting hosts from the source LUN to the clone LUN, should the source become corrupted. - Makes backup operation of a clone nondisruptive.	Provides a quick and instant copy because it is a virtual LUN (provided that the snapshot is used in conjunction with a backup operation).	Provides an extra level of protection against critical data loss if both the source LUN and clone LUN become corrupted.
Creation Time	Minutes to hours. The creation time depends on the size of the source LUN.	Instantaneous	Instantaneous
Disk Space Used	Uses the same amount of disk space as the source LUN.	Uses reserved LUN pool space, which is usually 10% to 20% of the source LUN size, but will vary depending on how much data has changed on the source LUN.	Uses reserved LUN pool space (for the snapshot) and full disk space (for the clone), which usually totals 110% to 120% of the source LUN size, but will vary depending on how much data has changed on the source LUN.
Disk Space Used	Uses the same amount of disk space as the source LUN.	Uses reserved LUN pool space, which is usually 10% to 20% of the source LUN size, but will vary depending on how much data has changed on the source LUN.	Uses reserved LUN pool space (for the snapshot) and full disk space (for the clone), which usually totals 110% to 120% of the source LUN size, but will vary depending on how much data has changed on the source LUN.
Data Recovery Time After Source LUN Failure/Corruption	Instantaneous after initializing a reverse synchronization.	Restore time from network or backup tape.	Restore time from network or backup tape if fractured clone received host I/O.

Table 4-2 A Comparison of Clones and Snapshots (*continued*)

	Clones	Snapshots	Snapshots of Clones
Performance Impact on the Storage System	- There is no performance impact when a clone is in a fractured state. - For the initial synchronization of the clone, performance decreases. Subsequent synchronizations have minimal impact on performance. - For synchronizations and reverse synchronizations, the performance decrease is the same as an initial synchronization but for a shorter period of time. Impact is also determined by the synchronization rate, which is set when the clone is added to the Clone Group.	A slight performance decrease due to the copy-on-first-write.	Combination of both clones and snapshots.

SnapView Worksheets

The following three worksheets will help you plan your clones and snapshots. The first worksheet is for clone planning; the second is for snapshot planning; the third is for reserved LUN pool planning for snapshots.

Clone Planning Worksheet

The following worksheet will help you plan for clones. Note that the size of the clone private LUNs (one for each SP) is fixed and it is quite modest, so the worksheet does not include it.

Clone Planning Worksheet

Production Host Name	LUN			
	Source LUN ID	Storage Group ID	Size (Gbytes)	Application, File System, or Database Name

Snapshot Planning Worksheet

For each SnapView session, you must complete a snapshot session worksheet. Note that you must assign the LUN and snapshot to different Storage Groups. One Storage Group must include the production host and source LUN; another Storage Group must include the second host and the snapshot.

Snapshot Session Worksheet

[illegible]

Since snapshots use the reserved LUN pool, you need the following information for system setup so you can bind LUNs for the reserved LUN pool. The number of reserved LUNs should always be twice the number of source LUNs. Note that MirrorView/A and incremental SAN Copy also use the reserved LUN pool. One application can use a reserved LUN as long as another application is not using it.

Reserved LUN Pool Setup Worksheet^a

SP (A or B)	Source LUN Size	RAID Type for Reserved LUN Pool	RAID Group ID of Parent RAID Group	Reserved LUN size ^b (Gbytes)	Reserved LUN ID (complete after binding)

- a. This worksheet is also provided in the MirrorView/A and incremental SAN Copy sections.
- b. EMC recommends 20% of the source LUN size.

What is MirrorView?

EMC MirrorView is storage-system-based software that maintains a copy of an image of a logical unit (LUN) at separate locations. The images are far enough apart to provide for disaster recovery, that is, to let one image continue if a serious accident or natural disaster disables the other.

The production image (the one mirrored) is called the primary image; the copy of the image is called the secondary image. The primary image is connected to a server called the production host. The secondary image is maintained by a separate storage system that can be a stand-alone storage system or connected to its own server. Since the same management station must manage both storage system, you can use it to promote the secondary image if the primary image becomes inaccessible.

An optional write intent log keeps track of writes that may not yet be recorded in one or more of a mirror's images. Using the write intent log provides the maximum protection against having to do a full resynchronization.

MirrorView also supports protocols that let you use FC/IP to send mirrored data over extended distances.

MirrorView Example

Figure 4-3 shows two sites and a primary and secondary image that includes the database of three LUNs.


Each server has a path to each SP on each storage system through each fabric. If a failure occurs in a path, then the storage-system software may switch to the path through the other SP (transparent to any applications).

The server sends a write request to an SP in Storage System 1, which then writes data to the local LUN and mirrors the data on on Storage System 2. The mirroring is synchronous, so that Storage System 2 always contains all data modifications that are acknowledged by Storage System 1 to the production host.

The server sends a write request to an SP in Storage System 1, which writes data to its image while sending it to the corresponding SP in Storage System 2. Storage System 2 then writes the data to its image.

Storage System 1 then acknowledges the write to the production host.

Database Server 2, the other host, has no direct access to the mirrored data. (The standby site does not need a server at all.) This server runs applications that access other data stored on Storage System 2. If a failure occurs in either the production host or Storage System 1, an administrator can use the management station to promote the image on Storage System 2 to the primary image. Then the appropriate applications can start on any connected server (here, Database Server 2) with full access to the data. The mirror will be accessible in minutes, although the time needed for applications to recover will vary.



EMC2000

Figure 4-3 Sample MirrorView Configuration

MirrorView Features and Benefits

MirrorView mirroring adds value to customer systems by offering the following features:

- ◆ Provision for disaster recovery with minimal overhead
- ◆ CLARiiON environment
- ◆ Bidirectional (cross) mirroring
- ◆ Integration with EMC SnapView LUN snapshot and clone copy software

CLARiiON Environment

MirrorView operates in a highly available environment, leveraging the dual-SP design of CLARiiON storage systems. If one SP fails, MirrorView running on the other SP controls and maintains the mirrored LUNs. If the server is able to fail over I/O to the remaining SP, then periodic updates continue. The highly-available features of RAID protect against disk failure, and mirrors are resilient to an SP failure in the primary of secondary storage system.

Provision for Disaster Recovery with Minimal Overhead

Provision for disaster recovery is the major benefit of MirrorView mirroring. Destruction of the primary data site would cripple or ruin many organizations. After a disaster, MirrorView lets data processing operations resume very quickly, often within minutes or a few hours.

MirrorView is transparent to servers and their applications. Server applications do not know that a LUN is mirrored, and the effect on performance depends on the distance and application workload.

MirrorView supports up to four connections per storage system, which means you can use it in a central backup configuration, where one storage system maintains remote mirrors for as many as four other storage systems.

MirrorView uses synchronous writes, which means that server writes are acknowledged only after all secondary storage systems commit the data. Most disaster recovery systems sold today use this type of mirroring.

MirrorView runs in storage systems, not on servers, therefore it uses no host I/O or CPU resources. The additional processing for mirroring is performed on storage systems.

Local High Availability

MirrorView operates in a highly available environment. Each host has two host bus adapters (HBAs) and each storage system has two SPs. If a single adapter or SP fails, the path in the surviving SP can take control of (*trespass*) any LUNs owned by the failed adapter or SP. The high-availability features of RAID protect against disk failure. Mirrors are resilient to an SP failure in the primary or secondary storage system.

Bidirectional Mirroring

The primary or secondary role applies to just one remote mirror. A storage system can maintain the primary image for one mirror and a secondary image for another mirror. This allows the use of server resources at both sites while maintaining duplicate copies of all data at both sites.

Integration with EMC SnapView Software

EMC SnapView software allows users to create up to eight snapshot copies of an active LUN at any point in time. Each snapshot copy is a consistent image that can serve as backup while I/O continues to the original LUN. You can use SnapView in conjunction with MirrorView to make backup copies at a remote site.

A common situation for disaster recovery is to have a primary and a secondary site that are geographically separate. MirrorView ensures that the data from the primary site replicates to the secondary site. The secondary site sits idle until there is a failure of the primary site. With the addition of SnapView at the secondary site, the secondary site can take snapshot copies of the replicated images and back them up to other media, providing time-of-day snapshots of data on the production host with minimal overhead.

MirrorView Planning Worksheet

The following two worksheets help you plan your remote mirrors. The first worksheet is for remote mirror planning, and the second is for the optional write intent log planning. You must assign each primary image LUN to a Storage Group (as with any normal LUN), but you must not assign a secondary image LUN to a Storage Group.

Remote Mirrors Worksheet

Production Host Name	Primary LUN ID, Size, and File System Name	Storage Group Number or Name	Use Write Intent Log - Y/N ^a	SP (A/B)	Remote Mirror Name	Secondary Host Name	Secondary Image LUN ID

- a. By default, newly created remote mirrors do not use the write intent log, so you do not need to allocate space for the write intent log. However, before you use it, you must allocate space for it. If you want to use the write intent log, answer Y, and, after completing the worksheet, fill out the Write Intent Log Worksheet on the next page.

If you decide to use the write intent log feature for even one remote mirror, then you must you must bind two LUNs of at least 128 Mbytes for the write intent log.

Optional Write Intent Log Worksheet

	LUN Name or ID	LUN Size (128 Mbytes minimum)
First LUN		
Second LUN		

What is MirrorView/A?

MirrorView / A is storage-system-based software for periodically updating remote copies of production data. It keeps a point-in-time copy of a logical unit (LUN) and periodically replicates the copy to a separate location to provide disaster recovery; that is, to let one image continue to be active if a serious accident or natural disaster disables the other. It provides data replication over long distances (hundreds to thousands of miles).

The production image (the one mirrored) is called the primary image; the copy image is called the secondary image. MirrorView / A supports one remote image per primary. The primary image receives I/O from a host called the production host; a separate storage system maintains the secondary image. This storage system can optionally have a failover/standby server connected to it or can be connected to its own server. Both storage systems are managed with the same domain. The management station can promote the secondary image if the primary image becomes inaccessible. After initial synchronization, the remote site always has a consistent copy of the primary data.

Like all remote mirroring products, MirrorView / A does not flush host buffers before replicating the primary data. Therefore, the copy is a crash-consistent image of the primary data. You must verify data integrity of the secondary image before using it for disaster recovery. (The verification process varies by application type.)

MirrorView / A supports MirrorView / A Consistency Groups, which this manual refers to as *Consistency Groups*. A Consistency Group is a collection of asynchronous mirrors that must remain intact during any disruptive situations. When you mirror different LUNs asynchronously that contain data with a content relationship to each other (such as RDBMS database files and transaction logs), the mirroring software must update the secondary images of these LUNs in a way that retains their point-in-time content relationship to each other. If this relationship is not retained, the application may not be able to use the data in the secondary image LUNs should a failure occur with a primary image LUN.

The LUNs in a MirrorView / A Consistency Group must reside on the same storage system, unlike the volumes in a Symmetrix Consistency Group, which can reside on multiple storage systems.

MirrorView/A Example

Figure 4-4 shows a sample remote mirror configuration with two sites and a primary and secondary image that includes a database of three LUNs.


Database Server 1, the production host, executes customer applications. These applications access data on Storage System 1, in the database server Storage Group. Storage System 2 mirrors the data on the Database Server Storage Group.

Storage Groups are different than Consistency Groups. A Storage Group is a group of LUNs assigned to a server. A server can access only the LUNs in the Storage Group to which it is assigned.

Each server has a path to each SP in each storage system through each fabric. If a failure occurs in a path, then the storage-system software may switch to the path through the other SP (transparent to any applications).

The server sends a write request to an SP in Storage System 1, which then writes data to the local LUN. The change to the primary LUN is recorded, and at an interval that you define, all changes are copied to the secondary storage system.

If a failure occurs in Storage System 1, an administrator can use the management station to promote the image on Storage System 2 to the role of primary image. Then the appropriate application can start on any connected server (here, Database Server 2) with full access to the data. The mirror will be accessible in minutes, although the time needed for applications to recover will vary. Database Server 2, the standby host, has no direct access to the mirrored data. (No server is necessary at the standby site; if there is no server, the storage-system SPs are still connected as shown.) This server can also run applications that access other data stored on Storage System 2.



EMC2000

Figure 4-4 Sample Remote Mirror Configuration

MirrorView/A Features and Benefits

MirrorView/A mirroring has the following features:

- ◆ Provision for disaster recovery with minimal overhead
- ◆ CLARiON environment
- ◆ Bidirectional mirroring
- ◆ Integration with EMC SnapView LUN copy software
- ◆ Replication over long distances

Provision for Disaster Recovery with Minimal Overhead

The provision for disaster recovery is the major benefit of MirrorView/A mirroring. Destruction of the data at the primary site would cripple or ruin many organizations. MirrorView/A lets data processing operations resume with minimal overhead. MirrorView/A enables a quicker recovery by creating and maintaining a copy of the data on another storage system.

MirrorView/A is transparent to hosts and their applications. Host applications do not know that a LUN is mirrored and the effect on performance is minimal.

With MirrorView/A, secondary systems are periodically updated.

MirrorView/A is not host-based; therefore it uses no host I/O or CPU resources. The processing for mirroring is performed on the storage system.

CLARiiON Environment

MirrorView/A operates in a highly available environment, leveraging the dual-SP design of CLARiiON storage systems. If one SP fails, MirrorView/A running on the other SP controls and maintains the mirrored LUNs. If the server is able to fail over I/O to the remaining SP, then periodic updates continue. The highly-available features of RAID protect against disk failure, and mirrors are resilient to an SP failure in the primary or secondary storage system.

Bidirectional Mirroring

A single storage system can be primary (that is, hold the primary image) for some mirrors and secondary (that is, hold the secondary image) for others. This enables bidirectional mirroring.

A storage system can never hold more than one image of a single mirror.

Integration with SnapView Software

SnapView software lets you create a snapshot of an active LUN at any point in time; however, do this only when the mirror is not updating the secondary image. The snapshot is a consistent image that can serve for other application purposes while I/O continues to the source LUN. The secondary image is not viewable to any hosts, but you can use SnapView in conjunction with MirrorView/A to create a snapshot of a secondary image on a secondary storage system to perform data verification and run parallel processes.

Before taking a snapshot, make sure that the secondary LUN is in a normal condition and not updating. A snapshot of a secondary LUN when MirrorView/A is updating the secondary will not give consistent data.

To provide for disaster recovery, the primary and secondary storage systems should be geographically separated. MirrorView/A ensures that data from the primary storage system replicates to the secondary. The server (if any) connected to the secondary might sit idle until the primary site fails. With SnapView at the secondary site, the server at the secondary site can take snapshots of the mirror images and back them up to other media. This provides point-in-time snapshots of production data with little impact to production host performance.

You can create a snapshot of a mirrored LUN, but you cannot create a clone (BCV) of a mirrored LUN.

Replication Over Long Distances

MirrorView/A uses FC-to-IP devices to provide replication over long distances (hundreds to thousands of miles).

MirrorView/A Planning Worksheets

The following two worksheets will help you plan your remote asynchronous mirrors. The first is for mirror planning, and the second is for reserved LUN pool planning.

You must assign each primary image LUN to a Storage Group (as with any normal LUN), but you must not assign a secondary image LUN to a Storage Group. If primary image LUNs are in a mirror group, all secondary image LUNs must be on the same remote storage system.

Remote Asynchronous Mirrors Worksheet

Production Host Name	Primary LUN ID, Size, and File System Name	Storage Group Number or Name	SP (A/B)	Remote Mirror Name	Secondary Host Name	Secondary Image LUN ID	Group Name (if applicable)

Since MirrorView /A uses the reserve LUN pool, you need the following information for system setup so you can bind LUNs for the reserved LUN pool. The number of reserved LUNs should always be twice the number of source LUNs. A MirrorView /A primary image LUN is a source LUN. Note that SnapView and incremental SAN Copy also use the reserved LUN pool. One application can use a reserved LUN as long as another application is not using it.

Reserved LUN Pool Setup Worksheet^a

SP (A or B)	Source LUN Size	RAID Type for Reserved LUN Pool	RAID Group ID of Parent RAID Group	Reserved LUN size ^b (Gbytes)	Reserved LUN ID or name (complete after binding)

- a. This worksheet is also provided in the SnapView and incremental SAN Copy sections.
- b. EMC recommends 20% of the source LUN size.

What Is EMC SAN Copy?

EMC SAN Copy is a storage-system-based software for copying data directly from a logical unit on one storage system to destination logical units on supported remote systems without using host resources. SAN Copy connects through a SAN, and also supports protocols that let you use the IP WAN to send data over extended distances (see *Implementing SAN Copy Over Extended Distances* on page 4-28).

SAN Copy runs in the SPs of a CX400, CX500, CX600, CX700, or FC4700 - Series storage system (called a SAN Copy storage system), and not on the host servers connected to the storage systems. As a result, the host processing resources are free for production applications while the SPs copy data.

SAN Copy can copy data between logical units as follows:

- ◆ Within a CLARiiON storage system
- ◆ On CLARiiON storage systems
- ◆ On CLARiiON and Symmetrix storage systems
- ◆ On CLARiiON and qualified non-EMC storage systems

SAN Copy can use any CLARiiON SP ports to copy data, provided the port is not being used for MirrorView connections. Multiple sessions can share the same port. You choose which ports SAN Copy sessions use through switch zoning.

You can use SAN copy to create full or incremental copies of a source logical unit. An incremental session copies only the data that has changed since the last copy session. This can significantly reduce the time needed to copy the data, thereby allowing the copy operation to be performed more frequently and more effectively. Unlike full copy sessions, the source logical unit for an incremental session can remain online during the copy process. Regardless of the type of copy session — full or incremental — the participating logical units must meet certain requirements.

For full and incremental sessions, the source and destination logical units must be

- ◆ the same size, or the destination logical unit must be larger.
- ◆ for use by the same operating system.
- ◆ owned by the same SP, for a copy within one storage system.

In addition, for full sessions:

- ◆ the source and destination logical units must be inactive (offline and/or unmounted) during the copy process.
- ◆ No writes to the source logical unit (unless you use replication software), and no application I/O with the destination logical unit can occur during the copy session

In addition, for incremental sessions, the source logical unit:

- ◆ *cannot* be a SnapView snapshot
- ◆ must reside on the SAN Copy storage system
- ◆ *can* remain online during the copy process.

You can use SAN Copy for the following tasks:

- ◆ One-time data migration
- ◆ Routine copying of data for application testing
- ◆ Content distribution
- ◆ Scheduled updates to multiple databases.

To manage SAN Copy sessions, you can use either the UI-based Navisphere Manager, or the SAN Copy Command Line Interface.

SAN Copy Example

Figure 4-5 shows two logical units being copied to two different storage systems. One SAN Copy session is copying LUN 15 in the CX600 storage system to volume 32 in the Symmetrix storage system. The second SAN Copy session is copying LUN 6 in the FC5300 storage system to LUN 14 in the CX600 storage system. In this example, SAN Copy is required only on the CX600.




Figure 4-5 Sample SAN Copy Sessions Between Storage Systems

SAN Copy Features and Benefits

SAN Copy adds value to customer systems by offering the following features:

- ◆ Storage-system-based data mover application that offloads the host, thereby improving host performance
- ◆ Software that you can use in conjunction with replication software, allowing I/O with the source logical unit to continue during the copy process
- ◆ Simultaneous sessions that copy data to multiple CLARiiON and Symmetrix storage systems
- ◆ Easy-to-use, web-based application for configuring and managing SAN Copy

Storage-System Based Data Mobility Application

SAN Copy is a storage-system based data mobility application that uses a SAN to copy data between storage systems. Since SAN Copy runs on the storage systems, the need to move data to and from the attached hosts is eliminated, freeing host resources for users and applications. Since the host is not involved in the copy process and the data migration takes place on the SAN, the copy process is much faster than a LAN-based, host-involved copy process.

Use with Replication Software

If the storage system containing the source logical unit has device replication software, such as EMC SnapView or EMC TimeFinder™, SAN Copy can use a snapshot, clone, or Symmetrix BCV (business continuation volume) as its source logical unit — allowing I/O with the source logical units for the snapshot, clone, or BCV source logical unit to continue during the copy process.

Simultaneous Sessions Between Multiple Storage Systems

SAN Copy lets you have more than one session active at the same time. The number of supported concurrent active sessions and the number of logical units per session depends on the storage-system type.

Implementing SAN Copy Over Extended Distances

SAN Copy supports protocols that let you use the IP WAN to send LUN data over extended distances. In order to do this, you need to extend your switch fabric by using Fibre Channel-to-IP converters that allow Fibre Channel devices to communicate over TCP/IP networks.

When using SAN Copy over IP, a read operation will be faster than a write operation.

Performance Considerations

To get the best performance when implementing SAN Copy over IP, it is important to consider the following performance factors:

- ◆ Network performance
- ◆ SAN Copy configuration
- ◆ Data compression

Network Performance

The quality of the network is a critical factor in determining SAN Copy performance, especially with respect to latency (delay in the network) and bandwidth (frequency of the transferred bits).


SAN Copy Configuration

In order to address the network performance issues, SAN Copy lets you configure the buffers the SAN Copy operations use.

Data Compressibility

Data compression allows more data to be sent in the same amount of time, effectively multiplying the effective bandwidth. The amount of data compression depends on many factors, such as data type, compression algorithm of the equipment, and line speed.

Figure 4-6 shows a sample SAN Copy over IP configuration. In this example, LUN data is being transferred over the IP storage network using two FC-IP converters. Depending on the direction of the data transfer, the FC-to-IP converter either converts Fibre Channel data to IP data or IP data to Fibre Channel data.



EMC2694

Figure 4-6 Sample SAN Copy Over IP Configuration

SAN Copy Planning Worksheet

The following three worksheets will help you plan SAN Copy sessions. The first worksheet is for sessions between storage systems; the second is for sessions within the same storage system; and the third is for reserved LUN pool planning.

One of the storage systems in the session must be a SAN Copy storage system; that is, it must have SAN Copy installed on it.

SAN Copy Worksheet for a Session Between Storage Systems

Session Name:
SAN Copy Storage System Name:
Source Storage-System Name: Source Logical Unit Name:
Destination Storage-System Name: Destination Logical Unit Names:

For a SAN Copy session between logical units in the same storage system, the same SP must own the logical units.

SAN Copy Worksheet for A Session Within Storage System

Session Name:
Source Storage-System Name:
Source Logical Unit Name:
Destination Logical Unit Names:

Incremental SAN Copy uses the reserve LUN pool, so if you want to use this type of SAN Copy, you need the following information for system setup so you can bind LUNs for the reserved LUN pool. The number of reserved LUNs should always be twice the number of source LUNs. Note that MirrorView/A and SnapView also use the reserved LUN pool. One application can use a reserved LUN as long as another application is not using it

Reserved LUN Pool Setup Worksheet^a

SP (A or B)	Source LUN Size	RAID Type for Reserved LUN Pool	RAID Group ID of Parent RAID Group	Reserved LUN size ^b (Gbytes)	Reserved LUN ID or name (complete after binding)

- a. This worksheet is also provided in the SnapView and MirrorView/A sections.
- b. EMC recommends 20% of the source LUN size.

What Next? To plan LUNs and file systems, go to Chapter 5.
For details on the storage-system hardware, go to Chapter 6.

Planning File Systems and LUNs

This chapter shows sample installations with sample shared switched and unshared direct storage, and then provides worksheets for planning your own storage installation.

Major sections are

- ◆ Multiple Paths to LUNs5-2
- ◆ Sample Shared Switched Installation.....5-2
- ◆ Sample Unshared Direct Installation5-5
- ◆ Planning Applications, LUNs, and Storage Groups.....5-6

Multiple Paths to LUNs

A shared storage-system installation includes two or more servers and one or more storage systems, each with two SPs and the Access Logix option. Often shared storage installations include two or more Fibre Channel switches.

Shared storage (switched or direct) provides at least two paths to each LUN in the storage system. The Access Logix software detects all paths and, using optional PowerPath® failover software, can automatically switch to the other path, without disrupting applications, if a device (such as a host bus adapter or cable) fails. And with two adapters and two SPs, PowerPath can send I/O to each available path in a user-selectable sequence (multipath I/O) for load sharing and greater throughput.

Unshared storage, if the server has two adapters, provides two paths to each LUN. With two adapters, PowerPath performs the same function as with shared systems: automatically switches to the other path if a device (such as host bus adapter or cable) fails.

Sample Shared Switched Installation

Figure 5-1 shows a sample shared storage system connected to three servers: two servers in a cluster and one server running a database management program.

Disk IDs have the form *b e d*, where

- b* is the back-end loop (also referred to as a back-end bus) number (0, or 1)
- e* is the enclosure number, set on the enclosure front panel (0 for the first — lowest — DAE2)
- d* is the disk position in the enclosure (left is 0, right is 14)

Note that Figure 5-1 shows a single-cabinet storage system. In a storage system with two cabinets, all the disks in the first cabinet are on loop 0 (so the disk IDs in the first cabinet begin with 0) and the disks in the second cabinet are on loop 1 (so their disk IDs begin with 1).




Figure 5-1 Sample Shared Switched Storage Configuration

The storage-system disk IDs, RAID types, storage types, capacities, and drive letters are as follows. The capacities shown reflect 73-Mbyte disks.

Database Server LUNs (DS) - SP A 1,440 Gbytes on Seven LUNs			Clustered System LUNs					
			Mail Server LUNs (MS) - SP A 1,460 Gbytes on Seven LUNs			File Server LUNs (FS) - SP B 1,460 Gbytes on Seven LUNs		
Disk IDs	RAID Type	Storage Type, Capacity, Drive	Disk IDs	RAID Type	Storage Type, Capacity, Drive	Disk IDs	RAID Type	Storage Type, Capacity, Drive
000-004	RAID 5	Dbase1, 292 Mbytes, drive G	010-014	RAID 5	Users, 292 Mbytes, drive N	020-024	RAID 5	UsersG_O, 292 Mbytes, drive T
005-009	RAID 5	Dbase2, 292 Mbytes, drive H	015-019	RAID 5	Specs, 292 Mbytes, drive O	025-029	RAID 5	UsersP_S, 292 Mbytes, drive U
0010-0014	RAID 5	Dbase3, 292 Mbytes, drive I	0110-0114	RAID 5	Apps, 292 Mbytes, drive P	0210-0214	RAID 5	UsersT_Z, 292 Mbytes, drive V
103-104	RAID 1	Log file,Dbase1, 73 Mbytes, drive J	110-114	RAID 5	ISP A Mail, 292 Mbytes, drive Q	120-124	RAID 5	Specs, 292 Mbytes, drive U
105-106	RAID 1	Log file, Dbase2, 73 Mbytes, drive K	115-119	RAID 5	ISP B Mail, 292 Mbytes, drive R	125-129	RAID 5	Apps, 292 Mbytes, drive V
107-108	RAID 1	Log file, Dbase3, 73 Mbytes drive L	1110-1114	RAID 5	ISP C Mail, 292 Mbytes, drive S	1210-1214	RAID 5	UsersA_F, 292 Mbytes, drive W
109-1014	RAID 5	Users, 345 Mbytes (6 disks), drive M						
100, 101, 102 – Hot spare (automatically replaces a failed disk in any server's LUN)								

Note: The available storage is 50% for a RAID 1 LUN and 80% for a 5-disk RAID 5 LUN.

Sample Unshared Direct Installation

This section shows the disks and LUNs in an unshared direct storage-system installation.




Figure 5-2 Unshared Direct Storage System

In the example below, each disk holds 73 Gbytes, so the storage system provides the server with 504 Gbytes of disk storage, all highly available. The storage-system disk IDs and LUNs are as follows:.

LUNs - SP A and SP B - 1,440 Gbytes on Seven LUNs		
Disk IDs	RAID Type	Storage type, capacity, drive
000-004	RAID 5	Dbase1, 292 Mbytes, drive G
005-009	RAID 5	Dbase2, 292 Mbytes, drive H
0010-0014	RAID 5	Mail, 292 Mbytes, drive I
101-102	RAID 1	Log file,Dbase1, 73 Mbytes, drive J
103-104	RAID 1	Log file, Dbase2, 73 Mbytes, drive K
105-109	RAID 1	Users, 292 Mbytes, drive L
1010-1014	RAID 5	Specs, 292 Mbytes, drive M
100 – Hot spare (automatically replaces a failed disk in any LUN)		

Note: The available storage is 50% for a RAID 1 LUN and 80% for a 5-disk RAID 5 LUN.

Planning Applications, LUNs, and Storage Groups

This section will help you plan your storage use — the applications to run, the LUNs that will hold them, and, for shared storage, the Storage Groups that will belong to each server.

Unshared storage systems do not use Storage Groups. For unshared storage, on the LUN and Storage Group worksheet, skip the Storage Group entry.

The worksheets include

- ◆ Application and LUN planning worksheet — lets you outline your storage needs.
- ◆ LUN and Storage Group planning worksheet — lets you decide on the disks to compose the LUNs and the LUNs to compose the Storage Groups for each server.
- ◆ LUN details worksheet — lets you plan each LUN in detail.

Make as many copies of each blank worksheet as you need. You will need this information later when you configure the storage system(s).

Sample worksheets appear later in this chapter.

Application and LUN Planning

Use the following worksheet to list the applications you will run, and the RAID type and size of LUN to hold them. For each application that will run, write the application name, file system (if any), RAID type, LUN ID (ascending integers, starting with 0), disk space required, and finally the name of the servers and operating systems that will use the LUN.

Application and LUN Planning Worksheet

Application	File System, Partition, or Drive	RAID Type of LUN	LUN ID (hex)	Disk Space Provided (Gbytes)	Server Hostname and Operating System

Sample Application and LUN Planning Worksheet

Application	File system, Partition, or Drive	RAID Type of LUN	LUN ID (hex)	Disk Space Provided (Gbytes)	Server Hostname and Operating System
<i>Dbase1</i>	<i>S</i>	<i>RAID 5</i>	<i>0</i>	<i>292 GB</i>	<i>Server1, Windows</i>
<i>Dbase2</i>	<i>T</i>	<i>RAID 5</i>	<i>1</i>	<i>292 GB</i>	<i>Server2, Windows</i>
<i>Log file for Dbase1</i>	<i>U</i>	<i>RAID 1</i>	<i>2</i>	<i>73 GB</i>	<i>Server1, Windows</i>
<i>Log file for Dbase2</i>	<i>V</i>	<i>RAID 1</i>	<i>3</i>	<i>73 GB</i>	<i>Server2, Windows</i>
<i>Users</i>	<i>W</i>	<i>RAID 5</i>	<i>4</i>	<i>365 GB</i>	<i>Server1, Windows</i>

Completing the Application and LUN Planning Worksheet

Application. Enter the application name or type.

File system, partition, or drive. Write the planned **file system**, **partition**, or **drive** name here.

With a Windows operating system, the LUNs are identified by drive letter only. The letter does not help you identify the disk configuration (such as RAID 5). We suggest that later, when you use the operating system to create a partition on a LUN, you use the disk administrator software to assign a volume label that describes the RAID type. For example, for drive T, assign the volume ID **RAID5_T**. The volume label will then identify the drive letter.

RAID type of LUN. This is the RAID Group type you want for this file system, partition, or drive. The features of RAID types are explained in Chapter 2. For a RAID 5, RAID 1, RAID 1/0, and RAID 0 Group, you can create one or more LUNs on the RAID Group. For other RAID types, you can create only one LUN per RAID Group.

LUN ID. The LUN ID is a hexadecimal number assigned when you bind the disks into a LUN. By default, the ID of the first LUN bound is 0, the second 1, and so on. Each LUN ID must be unique within the storage system, regardless of its Storage Group or RAID Group.

The maximum number of LUNs supported on one host bus adapter depends on the operating system.

Disk space required (Gbytes). Consider the largest amount of disk space this application will need, and then add a factor for growth.


If sometime in the future, you find that you need more space for this application, you can expand the capacity of the LUN by adding other LUNs to it. For information on expanding LUNs, refer to the Navisphere Manager online help or the *EMC Navisphere Manager Administrator's Guide*, P/N 069001125.

Server hostname and operating system. Enter the server hostname (or, if you don't know the name, a short description that identifies the server) and the operating system name, if you know it.

LUN and Storage Group Planning Worksheet

Use the appropriate worksheet to select the disks that will make up the LUNs and Storage Groups in each storage system. A storage system is any group of enclosures connected to a DPE2 (CX300 or CX500 storage system) or an SPE (CX700 storage system). A CX300 storage system can include up to three DAE2 enclosures for a total of 60 disks; a CX500 storage system can include up to seven DAE2 enclosures for a total of 120 disks; a CX700 storage system can include up to 16 DAE2 enclosures for a total of 240 disks. Unshared storage systems do not use Storage Groups. For unshared storage, skip the Storage Group entry.

LUN and Storage Group Worksheet for CX300



DAE2														
030	031	032	033	034	035	036	037	038	039	0310	0311	0312	0313	0314
DAE2														
020	021	022	023	024	025	026	027	028	029	0210	0211	0212	0213	0214
DAE2														
010	011	012	013	014	015	016	017	018	019	0110	0111	0112	0113	0114
DPE2														
000	001	002	003	004	005	006	007	008	009	0010	0011	0012	0013	0014
Vault disks														

Navisphere Manager displays disk IDs as *n-n-n*. The CLI recognizes disk IDs as *n_n_n*.

Storage-system number or name: _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN and Storage Group Worksheet for CX500

Loop 0 Enclosures

Loop 1 Enclosures

Navisphere Manager displays disk IDs as *n-n-n*. The CLI recognizes disk IDs as *n_n_n*.

DAE2															
130	131	132	133	134	135	136	137	138	139	1310	1311	1312	1313	1314	

DAE2															
030	031	032	033	034	035	036	037	038	039	0310	0311	0312	0313	0314	

DAE2															
120	121	122	123	124	125	126	127	128	129	1210	1211	1212	1213	1214	

DAE2															
020	021	022	023	024	025	026	027	028	029	0210	0211	0212	0213	0214	

DAE2															
110	111	112	113	114	115	116	117	118	119	1110	1111	1112	1113	1114	

DAE2															
010	011	012	013	014	015	016	017	018	019	0110	0111	0112	0113	0114	

DAE2															
100	101	102	103	104	105	106	107	108	109	1010	1011	1012	1013	1014	

DPE2															
000	001	002	003	004	005	006	007	008	009	0010	0011	0012	0013	0014	

Vault disks

Storage-system number or name: _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____


LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN and Storage Group Worksheet for CX700 - Single Cabinet



Storage-system number or name: _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____


LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN and Storage Group Worksheet for CX700 - Dual Cabinets



Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Storage Group ID or name: _____ Server hostname: _____


LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

LUN ID or name _____ RAID type _____ Cap. (Gb) _____ Disk IDs _____

Part of a sample LUN and Storage Group worksheet for one cabinet follows.



Storage-system number or name: SS1

Storage Group ID or name: 0 Server hostname: Server1 Dedicated ☒ Shared ☐

LUN ID or name 0 RAID type 5 Cap. (Gb) 292 Disk IDs 000, 001, 002, 003, 004

LUN ID or name 1 RAID type 5 Cap. (Gb) 292 Disk IDs 005, 006, 007, 008, 009

LUN ID or name 2 RAID type 5 Cap. (Gb) 292 Disk IDs 0010, 0011, 0012, 0013, 0015

LUN ID or name 3 RAID type 1 Cap. (Gb) 73 Disk IDs 101, 102

Storage Group ID or name: 0 Server hostname: Server1 Dedicated ☒ Shared ☐

LUN ID or name 4 RAID type 1 Cap. (Gb) 73 Disk IDs 103, 104

LUN ID or name 5 RAID type 5 Cap. (Gb) 292 Disk IDs 105, 106, 107, 108, 109

LUN ID or name 6 RAID type 5 Cap. (Gb) 292 Disk IDs 1010, 1011, 1012, 1013, 1014

LUN ID or name 7 RAID type 5 Cap. (Gb) 292 Disk IDs 1015, 1016, 1017, 1018, 1019

Disk Configuration Rules and Recommendations

The following rules and recommendations apply to CX300, CX500, and CX700 systems:

- ◆ You cannot use any of the disks 000 through 004 (enclosure 0, loop 0, disks 0-4) as a hot spare in a CX-Series system.
- ◆ The hardware reserves several gigabytes on each of disks 000 through 004 for the cache vault and internal tables. To conserve disk space, you should avoid binding any other disk into a RAID Group that includes any of these disks. Any disk you include in a RAID Group with a cache disk 000-004 is bound to match the lower unreserved capacity, resulting in lost storage of several gigabytes per disk.
- ◆ Each disk in the RAID Group should have the same capacity. All disks in a Group are bound to match the smallest capacity disk, and you could waste disk space. The first five drives (000-004) should always be the same size.

- ◆ You cannot mix ATA (Advanced Technology Attachment) and Fibre Channel disk drives within a RAID Group.
- ◆ Hot spares for Fibre Channel drives must be Fibre Channel drives; ATA drives require ATA hot spares.
- ◆ If a storage system will use disks of different speeds (for example, 10K and 15K rpm), then EMC recommends that you use disks of the same speed throughout each 15-disk enclosure. For any enclosure, the hardware allows *one* speed change within an enclosure, so if need be, you may use disks of differing speeds. Place the higher speed drives in the first (leftmost) drive slot(s).
- ◆ You should always use disks of the same speed and capacity in any RAID Group.
- ◆ Do not use ATA drives to store boot images of an operating system. You must boot host operating systems from a Fibre Channel drive.

Completing the LUN and Storage Group Planning Worksheet

As shown, draw circles around the disks that will compose each LUN, and within each circle specify the RAID type (for example, RAID 5) and LUN ID. This is information you will use to bind the disks into LUNs. For disk IDs, use the form shown on page 5-10.

Next, complete as many of the Storage System sections as needed for all the Storage Groups in the SAN (or as needed for all the LUNs with unshared storage). Copy the (blank) worksheet as needed.

For shared storage, if a Storage Group will be dedicated (not accessible by another server in a cluster), mark the Dedicated box at the end of its line; if the Storage Group will be accessible to one or more other servers in a cluster, write the hostnames of all servers and mark the Shared box.

For unshared storage, ignore the Dedicated/Shared boxes.

LUN Details Worksheet

Use the LUN Details Worksheet to plan the individual LUNs. Blank and sample completed LUN worksheets follow.

Complete as many blank worksheets as needed for all LUNs in storage systems. For unshared storage, skip the Storage Group entries.

LUN Details Worksheet

Storage system (complete this section once for each storage system)

Storage-system number or name: _____

Storage-system installation type

☐ Unshared Direct ☐ Shared-or-Clustered Direct ☐ Shared Switched

SP information: SP A: IP address or hostname: _____ Port ALPA ID: _____ Memory (Mbytes): _____

SP B: IP address or hostname: _____ Port ALPA ID: _____ Memory (Mbytes): _____

☐ Caching Read cache: _____% Write cache: _____% Cache page size (Kbytes): _____

☐ RAID-3

LUN ID: _____ SP owner: ☐ A ☐ B SP loop (0 or 1): _____

RAID Group ID: _____ Size, GB: _____ LUN size, GB: _____ Disk IDs: _____

RAID type: ☐ RAID 5 ☐ RAID 3 - Memory, MB: _____ ☐ RAID 1 mirrored pair ☐ RAID 0
☐ RAID 1/0 ☐ Individual disk ☐ Hot spare

Caching: ☐ Read and write ☐ Write ☐ Read ☐ None

Servers that can access this LUN's Storage Group: _____

Operating system information: Device name: _____ File system, partition, or drive: _____

LUN ID: _____ SP owner: ☐ A ☐ B SP loop (0 or 1): _____

RAID Group ID: _____ Size, GB: _____ LUN size, GB: _____ Disk IDs: _____

RAID type: ☐ RAID 5 ☐ RAID 3 - Memory, MB: _____ ☐ RAID 1 mirrored pair ☐ RAID 0
☐ RAID 1/0 ☐ Individual disk ☐ Hot spare

Caching: ☐ Read and write ☐ Write ☐ Read ☐ None

Servers that can access this LUN's Storage Group: _____

Operating system information: Device name: _____ File system, partition, or drive: _____

LUN ID: _____ SP owner: ☐ A ☐ B SP loop (0 or 1): _____

RAID Group ID: _____ Size, GB: _____ LUN size, GB: _____ Disk IDs: _____

RAID type: ☐ RAID 5 ☐ RAID 3 - Memory, MB: _____ ☐ RAID 1 mirrored pair ☐ RAID 0
☐ RAID 1/0 ☐ Individual disk ☐ Hot spare

Caching: ☐ Read and write ☐ Write ☐ Read ☐ None

Servers that can access this LUN's Storage Group: _____

Operating system information: Device name: _____ File system, partition, or drive: _____

LUN Details Worksheet

Storage system (complete this section once for each storage system)

Storage-system number or name: SS1

Storage-system installation type

☒ Unshared Direct ☐ Shared-or-Clustered Direct ☐ Shared Switched

SP information: SP A: IP address or hostname: SS1spa Port ALPA ID: 0 Memory(Mbytes): 1 Gb
 SP B: IP address or hostname: SS1spb Port ALPA ID: 1 Memory(Mbytes): 1 Gb

☐ Caching Read cache: 33% Write cache: 67% Cache page size (Kbytes): 4

☐ RAID-3

LUN ID: 0 SP owner: ☐A ☐B SP loop (0 or 1): 0

RAID Group ID: 0 Size, GB: 292 LUN size, GB: 292 Disk IDs: 000, 001, 002, 003, 004

RAID type: ☒ RAID 5 ☐ RAID 3 - Memory, MB: ☐ RAID 1 mirrored pair ☐ RAID 0
☐ RAID 1/0 ☐ Individual disk ☐ Hot spare

Caching: ☐ Read and write ☐ Write ☐ Read ☐ None

Servers that can access this LUN's Storage Group: Server1

Operating system information: Device name: File system, partition, or drive: S

LUN ID: 1 SP owner: ☐A ☐B SP loop (0 or 1):

RAID Group ID: 1 Size, GB: 292 LUN size, GB: 292 Disk IDs: 005, 006, 007, 008, 009

RAID type: ☒ RAID 5 ☐ RAID 3 - Memory, MB: ☐ RAID 1 mirrored pair ☐ RAID 0
☐ RAID 1/0 ☐ Individual disk ☐ Hot spare

Caching: ☐ Read and write ☐ Write ☐ Read ☐ None

Servers that can access this LUN's Storage Group: Server1

Operating system information: Device name: File system, partition, or drive: T

LUN ID: 2 SP owner: ☐A ☐B SP loop (0 or 1): 1

RAID Group ID: 2 Size, GB: 73 LUN size, GB: 73 Disk IDs: 0101, 1011, 1012, 1013, 1014

RAID type: ☐ RAID 5 ☐ RAID 3 - Memory, MB: ☒ RAID 1 mirrored pair ☐ RAID 0
☐ RAID 1/0 ☐ Individual disk ☐ Hot spare

Caching: ☐ Read and write ☐ Write ☐ Read ☐ None

Servers that can access this LUN's Storage Group: Server1

Operating system information: Device name: File system, partition, or drive: U

Completing the LUN Details Worksheet

Complete the header portion of the worksheet for each storage system as described next. Copy the blank worksheet as needed.

Storage-System Entries

Storage-system installation type: specify Unshared Direct, Shared-or-Clustered Direct, or Shared Switched.

SP information: IP address or hostname. The IP address is required for communication with the SP. You don't need to complete it now, but you will need it when the storage system is installed so that you can set up communication with the SP.

Port ALPA ID. For shared storage (switches), the SP ALPA ID is ignored for all except HP/UX servers. For unshared storage, each SP connected to a server must have a unique SP Port ALPA ID; for example, for two storage systems, use IDs 0, 1, 2, and 3. The SP Port ALPA ID, like the IP address, is generally set at installation.

Memory (Mbytes). Each CX300 has a minimum of 1 Gbyte and maximum of 2 Gbytes of memory. Each CX500 SP has a minimum of 2 Gbytes and maximum of 4 Gbytes of memory. Each CX700 SP has a minimum of 4 Gbytes and maximum of 8 Gbytes of memory.

Caching. You can use SP memory for read/write caching or RAID 3. (We do not recommend using both caching and RAID 3 in the same storage system.) You can use different cache settings for different times of day. For example, for user I/O during the day, use more write cache; for sequential batch jobs at night, use more read cache. You enable caching for specific LUNs — allowing you to tailor your cache resources according to priority. If you choose caching, check the box and continue to the next cache item; for RAID 3, skip to entry **RAID 3**.

Read cache percentage. If you want a read cache, generally it should be about one third (33%) of the total available cache memory.

Write cache percentage. Generally, the write cache should be two thirds (67%) of the total available cache memory.

Cache page size. This applies to both read and write caches. It can be 2, 4, 8, or 16 Kbytes.

As a general guideline, we suggest 8 Kbytes. The ideal cache page size depends on the operating system and application.

RAID 3. If you want to use the SP memory for RAID 3, check the box.

RAID Group/LUN Entries

Complete a RAID Group/LUN entry for each LUN and hot spare.

LUN ID. The LUN ID is a hexadecimal number assigned when you bind the disks into a LUN. By default, the ID of the first LUN bound is 0, the second 1, and so on. Each LUN ID must be unique within the storage system, regardless of its Storage Group or RAID Group.

The maximum number of LUNs supported on one host bus adapter depends on the operating system.

SP owner. Specify the SP that will own the LUN: SP A or SP B. You can let the management program automatically select the SP to balance the workload between SPs; to do so, leave this entry blank.

SP back-end loops (buses). Each CX300 SP supports one physical loop that is paired with its counterpart SP to create a redundant loop (loop 0).

Each CX500 SP supports two physical loops that are paired with its counterpart SP to create four redundant loops (loop 0 and 1).

Each CX700 SP supports four physical loops that are paired with its counterpart SP to create eight redundant loops (loop 0, 1, 2, and 3).

Ideally, you will place the same amount of load on each loop. This may mean placing two or three heavily used LUNs on one loop, and six or eight lightly used LUNs on the other loop. The loop designation appears in the disk ID (form *b-e-d*, where *b* is the back-end loop number, *e* is the enclosure number, and *d* is the disk position in the enclosure. For example, 0-1-3 indicates the disk on loop 0, in enclosure 1, in the third position (fourth from left) in the storage system.

RAID Group ID. This ID is a hexadecimal number assigned when you create the RAID Group. By default, the number of the first RAID Group in a storage system is 0, the second 1, and so on, up to the maximum of 1F (31).

Size (RAID Group size). Enter the user-available capacity in gigabytes (Gbytes) of the whole RAID Group. You can determine the capacity as follows:

RAID 5 or RAID 3 Group: $\text{disk-size} * (\text{number-of-disks} - 1)$

RAID 1/0 or RAID 1 Group: $(\text{disk-size} * \text{number-of-disks}) / 2$

RAID 0 Group: $\text{disk-size} * \text{number-of-disks}$

Individual unit: disk-size

For example,

- A five-disk RAID 5 or RAID 3 Group of 73-Gbyte disks holds 292 Gbytes;
- An eight-disk RAID 1/0 Group of 73-Gbyte disks also holds 292 Gbytes;
- A RAID 1 mirrored pair of 73-Gbyte disks holds 73 Gbytes; and
- An individual disk of an 73-Gbyte disk also holds 73 Gbytes.

LUN size. Enter the user-available capacity in gigabytes (Gbytes) of the LUN. You can make this the same size as the RAID Group, described previously. Or, for a RAID 5, RAID 1, RAID 1/0, or RAID 0 Group, you can make the LUN smaller than the RAID Group. You might do this if you wanted a RAID 5 Group with a large capacity and wanted to place many smaller capacity LUNs on it; for example, to specify a LUN for each user. However, having multiple LUNs per RAID Group may adversely impact performance. If you want multiple LUNs per RAID Group, then use a RAID Group/LUN series of entries for each LUN.

Disk IDs. Enter the IDs of all disks that will make up the LUN or hot spare. These are the same disk IDs you specified on the previous worksheet. For example, for a RAID 5 Group in DAE2 on loop 0 (disks 10 through 14), enter **0010, 0011, 0012, 0013, and 0014.**

RAID type. Copy the RAID type from the previous worksheet. For example, RAID 5 or hot spare. For a hot spare (not strictly speaking a LUN at all), skip the rest of this LUN entry and continue to the next LUN entry (if any).

If this is a RAID 3 Group, specify the amount of SP memory for that group. To work efficiently, each RAID 3 Group needs at least 6 Mbytes of memory.

Caching. If you want to use caching (entry on page 5-17), you can specify the type of caching you want — read and write, read, or write for this LUN. Generally, write caching improves performance far more than read caching. The ability to specify caching on a LUN basis provides additional flexibility, since you can use caching for only the units that will benefit from it. Read and write caching recommendations follow.

Table 5-1 Cache Recommendations for Different RAID Types

RAID 5	RAID 3	RAID 1	RAID 1/0	RAID 0	Individual Unit
Highly Recommended	Not allowed	Recommended	Recommended	Recommended	Recommended

Servers that can access this LUN’s Storage Group. For shared switched storage or shared-or-clustered direct storage, enter the name of each server (copied from the LUN and Storage Group worksheet). For unshared direct storage, this entry does not apply.

Operating system information: Device name. Enter the operating system device name, if this is important and if you know it. Depending on your operating system, you may not be able to complete this field now.

File system, partition, or drive. Write the name of the file system, partition, or drive letter you will create on this LUN. This is the same name you wrote on the application worksheet.

On the following line, write any pertinent notes; for example, the file system mount- or graft-point directory pathname (from the root directory). If any of this storage system’s LUNs will be shared with another server, and the other server is the primary owner of this LUN, write **secondary**. (As mentioned earlier, if the storage system will be used by two servers, we suggest you complete one of these worksheets for each server.)

What Next? If you have completed the worksheets to your satisfaction, you are ready to learn about the hardware needed for these systems as explained in Chapter 6.

Storage-System Hardware

This chapter describes the storage-system hardware components.

Major sections are

- ◆ Storage-System Installation Types.....6-2
- ◆ Hardware for CX300, CX500, and CX700 Storage Systems6-3
- ◆ Planning Your Hardware Components6-8
- ◆ Hardware Dimensions and Requirements6-9
- ◆ Cabinets for Enclosures.....6-10
- ◆ Data Cable and Configuration Guidelines6-10
- ◆ Hardware Planning Worksheets6-12


Storage-System Installation Types

The storage systems attach to the server and the interconnect components as described in Chapter 1. The installation types are

- ◆ **Unshared direct**, with one server, is the simplest and least costly;
- ◆ **Shared-or-clustered direct**, with a limit of two servers, lets two servers share storage resources with high availability; and
- ◆ **Shared switched**, which has two switch fabrics, lets two or more servers share the resources of several storage systems in a storage area network (SAN).

Table 6-1 illustrates the storage-system installation types.

CX300, CX500, and CX700 storage systems require at least one network connection for management.



EMC1826

Figure 6-1 Types of Storage-System Installation

Hardware for CX300, CX500, and CX700 Storage Systems

The main hardware component for a CX300 or CX500 storage system is a disk processor enclosure (DPE2); for a CX700 storage system, it is a storage processor enclosure (SPE). Each storage system has two storage processors (SPs).

Table 1-1 on page 1-2 outlines each storage-system enclosure and disk capacity.

Storage Hardware

A CX300 or CX500 DPE2 or a CX700 SPE enclosure is a sheet-metal housing with a front door, midplane, and slots for the two storage processors (SPs), two power supplies, and fan modules. A CX300 and a CX500 DPE2 also has space for 15 disks. All components are replaceable units (FRUs) that can be replaced under power.

A separate standby power supply (SPS) is required to support write caching. Two are required for the CX500 and CX700; only one SPS is required for CX300, but it is more highly available with two.

Figure 6-2 shows the SPE, DAE2s, SPSs, and cabinet components of a CX700 storage system. Components of a CX300 and CX500 storage systems are similar.




Figure 6-2 CX700 Storage System

Storage Processor (SP)

The SP provides the intelligence of the storage system. Using its own proprietary software (called FLARE™ Operating Environment), the SP processes the data written to or read from the disk modules, and monitors the modules themselves. A CX700 SP consists of a printed-circuit board with memory modules, a board that provides SAN features (called a personality board), and status lights.

Each SP has front-end ports to communicate with switches or servers and back-end ports to communicate with disks. A CX300 DPE2 SP has two front-end ports and one back-end port; a CX500 DPE2 SP has two front-end ports and two back-end ports; a CX700 SPE SP has four front-end ports and four back-end ports.

For high availability, a storage system comes with two SPs. The second SP provides a second route to a storage system and also lets the storage system use write caching (below) for enhanced write performance.




Figure 6-3 Shared Storage Systems

Chapter 5 shows more examples of storage configurations.

Storage-System Caching

Storage-system caching improves read and write performance for several types of RAID Groups. Write caching, particularly, helps write performance — an inherent problem for RAID types that require writing to multiple disks. Read and write caching improve performance in two ways:

- ◆ For a read request — If a read request seeks information that's already in the read or write cache, the storage system can deliver it immediately, much faster than a disk access can.
- ◆ For a write request — the storage system writes updated information to SP write-cache memory instead of to disk, allowing the server to continue as if the write had actually completed. The write to disk from cache occurs later, at the most expedient time. If the request modifies information that's in the cache waiting to be written to disk, the storage system updates the information in the cache before writing it to disk; this requires just one disk access instead of two.

A standby power supply (SPS) protects data in the cache from power loss. If line power fails, the SPS provides power to let the storage system write cache contents to the vault disks. The vault disks are standard disk modules that store user data but have space reserved outside operating system control. When power returns, the storage system reads the cache information from the vault disks, and then writes it to the file systems on the disks. This design ensures that all write-cached information reaches its destination.

During normal operation, no I/O occurs with the vault; therefore, a disk's role as a vault disk has no effect on its performance.

SP Network Connection

Each SP has an Ethernet connection through which the EMC Navisphere® management software lets you configure and reconfigure the LUNs and Storage Groups in the storage system. Each SP connects to a network; this lets you reconfigure your system, if needed, should one SP fail.


EMC Navisphere is a member of the EMC ControlCenter™ family of products.

Disks

The disks — available in differing capacities — fit into slots in the DAE2 enclosure. EMC offers two types of DAE2:

- ◆ a standard DAE2 version with high-performance Fibre Channel disks and FC-AL link control cards
- ◆ a DAE2-ATA version with economical ATA (Advanced Technology Attachment) disks and FC-to-ATA link control cards to manage the disks and provide an interface between the FC-AL and the drives.

Each disk has a unique ID that you use when binding or monitoring its operation. The ID is derived from the Fibre Channel loop number, enclosure address and disk slot in the enclosure.



* Disks in DPE2 for CX300 or CX500 and in standard DAE2 for CX700

EMC2786

Figure 6-4 Disks and Their IDs

Disks with IDs 000 through 004 must be Fibre Channel disks. For a CX300 or a CX500 storage system, the slots for these disks are in the DPE2. For a CX700 storage system, the slots for these disks must be in a standard DAE2 and *not* a DAE2-ATA.

Enclosure 0 on loop 0 in a CX300, CX500, or CX700 storage system must contain disks with IDs 000 through 004. The remaining disk slots can be empty.

Planning Your Hardware Components

This section helps you plan the hardware components — adapters, cables, and storage systems and site requirements — for each server in your installation.

For shared switched storage or shared-or-clustered direct storage, you must use high-availability options: two SPs per storage system and at least two HBAs per server. For shared switched storage, two switch fabrics are required.

For unshared direct storage, a server may have one or two HBAs.

Components for Shared Switched Storage

The minimum hardware configuration required for shared switched storage is two servers, each with two host bus adapters, two Fibre Channel switch fabrics with one switch per fabric, and two SPs per storage system. Two SPS units (standby power supplies) are also required. You can use more servers (up to 15 are allowed), more switches per fabric, and more storage systems (up to four are allowed).

Components for Shared-or-Clustered Direct Storage

The minimum hardware required for shared-or-clustered direct storage is two servers, each with two host bus adapters, and one storage system with two SPs. You can use more storage systems (up to four are allowed).

Components for Unshared Direct Storage

The minimum hardware required for unshared direct storage is one server with two host bus adapters and one storage system with two SPs. You can choose more storage systems (up to four are allowed).

Hardware Dimensions and Requirements

The DPE2, SPE and SPS units, DAE2s, and switches require the following amounts of space and power in a standard 19-inch cabinet.

Table 6-1 CX300, CX500, and CX700 Hardware Dimensions and Requirements

Device	Dimensions	Vertical Size	Weight	Power	VA Rating
DPE2	13.34 cm (5.25 in) x 45.01 cm (17.72 in) x 60.33 cm (23.75 in)	3U (1.75 in per 1U)	44.5 kg (98.3 lb) with SPS, 15 disks, and mounting rails.	100 V-240 V ac At 100 V ac, 5.9 A. At 200 V ac, 2.9 A.	590 (fully configured) ^a
SPE	17.78 cm (7.0 in) x 44.45 cm (17.5 in) x 67.11 cm (26.42 in)	4U (1.75 in per 1U)	32.9 kg (72.6 lb) with rails	100 V-240 V ac At 100 V ac, 5.2 A max At 200 V ac, 2.6 A max	520 (fully configured) ^a
DAE2 (FC)	13.34 cm (5.25 in) x 45.01 cm (17.72 in) x 60.33 cm (23.75 in)	3U (1.75 in per 1U)	43.2 kg (95.3 lb) with LCCs, 15 disks, and mounting rails	100 V-240 V ac At 100 V ac, 4.0 A. At 200 V ac, 1.6 A	400 max (fully configured) ^a
DAE2-ATA				100 V-240 V ac At 100 V ac, 3.0 A. At 200 V ac, 1.4 A	300 max (fully configured) ^a
SPS Tray	4.01 cm (1.58 in) x 20.96 cm (8.25 in) x 60.33 cm (23.75 in)	1U (1.75 in per 1U)	25.9 kg (57.15 lb) with tray (two SPS units at 25 lb each)	100 V-240 V ac At 100 V ac, 0.2 A max At 200 V ac, 0.1 A max	60 max

a. A fully configured DAE2 includes 15 disk drives

The total of the device VA ratings cannot exceed 4800 VA per cabinet.

Cabinets for Enclosures

Pre-wired 19-inch-wide cabinets, ready for installation, are available in the following dimensions to accept storage-system enclosures.

Dimension	40U-C Cabinet	40U Cabinet*
Height (internal, usable for devices)	40U (179 cm; 70 inches) from floor pan to cabinet top (fan installed)	
Height (overall)	190 cm (75 inches)	
Width (usable for devices)	NEMA 19-inch standard; rail holes 45.78 cm (18.312 inches) apart center-to-center	
Width (overall)	60 cm (24 inches)	
Depth (front to back rail)	60 cm (24 inches)	
Depth (overall)	98.425 cm (39.37 inches); 103.75 cm (41.5 inches) with optional front door	91 cm (36 inches); 101 cm (40 inches) with optional front door
Weight (empty)	390 lb (177 kg) maximum without front door 440 lb (200 kg) maximum with front door	164.7 kg (363 pounds)
Maximum total device weight supported	945 kg (2100 pounds)	

*The 40U-C cabinet replaced 40U models in March, 2004.

Both cabinets accept combinations of:

- ◆ SPEs at 4U,
- ◆ DPE2s at 3U
- ◆ DAE2s at 3U,
- ◆ SPS units at 1U,
- ◆ Switches at 1 or 2U

The cabinets require 200–240 volts ac at 50/60 Hz, and include two to four power strips with compatible outlets. Plug options are L6–30P and IEC 309 30 A.

Filler panels of various sizes are available.

Data Cable and Configuration Guidelines

CX300, CX500, and CX700 storage systems require optical cable between servers, switches, and SPs. The cabling between SPE or DPE2 and DAE2 enclosures is copper.

EMC standard multimode optical configurations use 50 micron cable. EMC does not recommend mixing 62.5 micron and 50 micron optical cable in the same link. In certain situations, however, you can add a 50 micron adapter cable to the end of an already installed 62.5 micron cable plant. Contact your customer service representative for details.

Generally, you should minimize the number of cable connections, since each connection degrades the signal slightly and shortens the maximum distance the signal can carry.

Table 6-2 Data Cable Types and Sizes


Length	Typical Use
1 m (3.3 ft) copper	In a single cabinet, connects the SPE or DPE2 to a DAE2. In the same cabinet, connects a DAE2 to a DAE2.
5 m (16.5 ft) equalized copper	Cabinet to cabinet, connects the SPE or DPE2 to a DAE2.
10 m (33 ft) equalized copper	Cabinet to cabinet, connects a DAE2 to a DAE2.
3 m (10 ft) optical, LC-LC	In a single cabinet, connects the SPE or DPE2 to switches.
10 m (33 ft) optical, LC-LC 30 m (100 ft) optical LC-LC 50 m (164 ft) optical, LC-LC	In a room, connects the SPE or DPE2 to switches or connects the SPE to an HBA. (Adapter must support optical cable.)
100 m (328 ft) optical, LC-LC	In a room, connects the SPE to switches or connects the SPE to an HBA. (Adapter must support optical cable.)
300 m (1000 ft, 0.2 mi) optical, LC-LC	In a room, connects SPE to switches or connects the SPE to an HBA. (Adapter must support optical cable.)
10 m (33 ft) optical, LC-SC 30 m (100 ft) optical LC-SC 50 m (164 ft) optical, LC-SC 100 m (328 ft) optical, LC-SC 300 m (1000 ft, .2 mi) optical, LC-SC	LC-SC cables are used on the EMC 1-Gbit Fibre Channel products. You can use LC-SC cables to connect an HBA to a switch, a switch to the SPE, or an HBA to the SPE.
The minimum bend radius for optical cables is 3 cm (1.2 in). For the 50-micron cable operating at 2.125 Gbit/s the recommended maximum distance is 300 meters and at 1.0625 Gbit/s, the recommended, maximum distance is 500 meters. For the 62.5-micron cable operating at 2.125 Gbit/s the recommended maximum distance is 150 meters and at 1.0625 Gbit/s, the recommended maximum distance is 300 meters. All of these maximum lengths assume no more than two connections/splices in the link.	

Component planning diagrams and worksheets follow.

Hardware Planning Worksheets

Following are worksheets to note hardware components you want. Some installation types do not have switches or multiple servers.

Cable Planning Template for a Shared Storage System



EMC2449

Note: The cable identifiers apply to all storage systems.

Hardware Component Worksheet

Number of servers:____ Adapters in servers:____ Switches: 32-port:____ 24-port:____ 16-port:____ 8-port:____

SPEs or DPE2s:____ Rackmount cabinets:____

DAE2s:____LCCs:____

Cables between server and switch port- Cable A

Cable A₁, Optical: Number:____ Length____m or ft

Cable A₂, Optical: Number:____ Length____m or ft

Cable A_n, Optical: Number:____ Length____m or ft

Cables between switch ports and storage-system SP ports - Cable D


Cable D₁, CX300/500, up to 2/SP; CX700, up to 4/SP, Optical: Number:____ Length____m or ft

Cable D_m, CX300/500, up to 2/SP; CX700, up to 4/SP, Optical: Number:____ Length____m or ft

Cables between enclosures - Cable F

Cable F₂ (Copper, 1.0 m): Number (2 per DAE2): _____

Sample Cable Diagram for a Shared Storage System



EMC2436

Sample Hardware Component Worksheet

Servers: 3 Adapters in servers: 6 Switches: 32-port: _____ 24-port: _____ 16-port: 2 8-port: _____

SPEs or DPE2s: 1 Rackmount cabinets: 1

DAE2s: 6 LCCs: 12

Cables between server and switch - Cable A

Cable A₁, Optical: Number: 4 Length 10 m or ft

Cable A₂, Optical: Number: 2 Length 33 m or ft

Cable A_n, Optical: Number: _____ Length _____ m or ft

Cables between switch ports and storage-system SP ports - Cable D

Cable D₁ CX300/500, up to 2/SP; CX700, up to 4/SP, Optical: Number: 2 Length 33 m or ft

Cable D_m CX300/500, up to 2/SP; CX700, up to 4/SP, Optical: Number: _____ Length _____ m or ft

Cables between enclosures

Cable F₂ (Copper, 1.0 m): Number (2 per DAE): 12

What Next? If you have completed the worksheets to your satisfaction, then you are ready to order some of this equipment.

A

- ac power requirements, storage system 6-10
- Access Logix option 1-10
- ALPA ID 5-17
- application
 - planning
 - LUN 5-7
 - MirrorView 4-15, 4-22
 - SAN Copy 4-30
 - SnapView 4-8
 - sample, for RAID groups 2-20
- array, *see* disk-array storage system
- ATA (Advanced Technology Attachment), *see* DAE2 (2-gigabit disk-array enclosure)
- ATA disks, configuration rules 5-13

B

- backup, benefit with SnapView 4-3
- benefits
 - clones and snapshots 4-8
 - MirrorView 4-13
 - SAN Copy 4-27
 - SnapView 4-7
- bidirectional mirroring
 - MirrorView 4-14
 - MirrorView/A 4-20
- business continuance volumes (BCVs), *see* clone

C

- cabinets for rackmount storage systems 6-10

cabling

- guidelines 6-10
- types and sizes 6-11

cache, storage-system

- size 5-17

capacity

- disk 5-13
- RAID Groups and LUNs 5-19

cascading switches

- 1-6

chunk (snapshot) size

- 4-9

clone

- compared to snapshot 4-7
- definition of
- ID 4-3
- overview

components, storage system

- 6-3

configurations

- LUN and file system planning 5-6
- RAID, planning 5-6
- see also* installation types

Copy-On-First-Write, definition of

- 4-5

cross mirroring

- MirrorView 4-14
- MirrorView/A 4-20

D

- DAE2 (2-gigabit disk-array enclosure)
 - ATA 1-2, 6-7
 - dimensions and weight 6-9
 - standard 1-2, 6-7
 - types of 1-2, 6-7

- data mobility application 4-27
 - see also* SAN Copy
- data replication application 4-19
 - see also* SnapView, MirrorView, or MirrorView/A
- data sheets, hardware 6-10
- decision support
 - benefit with SnapView 4-3
- device name, operating system 5-20
- disaster recovery
 - MirrorView/A 4-20
- disaster recovery benefit
 - MirrorView 4-13
- disk
 - capacity 5-13
 - configuration rules 5-13
 - LUNs, planning 5-6
 - mirror, defined 2-2
 - module IDs 5-4
 - RAID types
 - compared 2-13
 - guidelines 2-18
 - selecting for RAID Groups 5-13
 - shared storage systems, examples 5-2
 - speeds 5-13
 - unit number on worksheet 5-9, 5-18
- disk-array storage system
 - application and LUN planning worksheet 5-7
 - cabinets 6-10
 - components 6-3
 - dimensions 6-9
 - hardware requirements shared/unshared 6-8
 - installation types 1-9
 - remote mirror example
 - MirrorView 4-11
 - MirrorView/A 4-18
 - sample shared installation 5-2
 - site requirements 6-9, 6-10
 - SP 6-5
 - unshared 5-5
- DPE2 (Disk Processor Enclosure, 2-Gbit)
 - dimensions and weight 6-9
 - statistics 6-9
- dual paths to LUNs 5-2

E

- enclosure address (EA) 6-7

F

- fabric, switch, introduced 1-4
- Fibre Channel
 - adapter 1-5
 - components 1-5
 - defined 1-4
 - switch, description 1-6
 - see also* switch
- file system
 - name 5-20
 - planning 5-6
- footprint
 - SPE 6-9
 - storage system 6-10
- FRUs (replaceable units) 6-3

G

- global spare, *see* hot spare

H

- hardware
 - data sheets 6-10
 - mirroring 2-2
 - mirroring, *see* MirrorView or MirrorView/A
 - planning 6-10
 - planning worksheets 6-12
 - storage 6-3
 - trade-offs 6-8
- height, storage system 6-10
- high availability
 - hardware trade-offs 6-8
 - remote mirror
 - MirrorView 4-14
 - shared switched installation 1-9
- host, *see* server
- hot spare
 - defined 2-11
 - restrictions 5-14
 - sample applications 2-21
 - when to use 2-19

I

- image (mirror)
 - defined 2-2
 - see also* MirrorView or MirrorView / A
- individual access array, *see* RAID 5 Group
- individual disk unit
 - defined 2-10
 - performance 2-15
 - sample applications 2-21
 - when to use 2-19
- installation
 - LUN and file system planning 5-6
 - sample 5-2
 - types 1-9
 - sample unshared direct 5-5
 - sample, shared switched 5-2
- interconnect components, cables, hubs, switches 1-5
- IP address of SP 5-17

L

- logical volume, *see* file system
- LUN (logical unit)
 - capacity 5-19
 - configurations, planning 5-6
 - disk mirror, defined 2-2
 - hot spare 2-11
 - in RAID Group 2-3
 - individual disk 2-10
 - number on worksheet 5-8, 5-18
 - paths to 5-2
 - planning file systems 5-6
 - planning worksheet 5-7, 5-14
 - RAID 0 Group 2-10
 - RAID 1 mirrored pair 2-7
 - RAID 1/0 Group 2-8
 - RAID 5 Group 2-4
 - RAID types
 - compared 2-5, 2-13
 - LUN (logical unit), Chapter 2
 - shared storage examples 5-2
 - unshared storage example 5-5
 - worksheet 5-15

M

- memory, SP, specifying on worksheet 5-17
- mirrored pair, *see* RAID 1 mirrored pair
- mirrored RAID 0 Group, *see* RAID 1/0 Group
- mirroring
 - defined 2-2
 - remote, *see* MirrorView software or remote mirroring
- MirrorView
 - benefits 4-13
 - bidirectional mirroring 4-14
 - disaster recovery 4-13
 - example 4-11
 - high availability 4-14
 - integration with SnapView 4-14
 - planning worksheet 4-15
 - primary image 4-11
 - remote mirroring, promotion 4-12
 - secondary image, promotion 4-12
- MirrorView / A
 - benefits
 - bidirectional mirroring 4-20
 - disaster recovery 4-20
 - example 4-18
 - integration with SnapView 4-20
 - planning worksheets 4-22
 - primary image 4-17
 - remote mirroring, promotion 4-18
 - replication over long distances benefit 4-21
 - secondary image, promotion 4-18
- multipath I/O 1-9
- multiple paths to LUNs 5-2

N

- Navisphere 3-3
- node, defined 1-4
- nonredundant array, *see* RAID 0 Group

O

- operating system
 - device name for disk unit 5-20
 - software mirroring 2-2
- optical cable, types and sizes 6-11

P

- page size, storage-system cache 5-17
- parallel access array, *see* RAID 3 Group
- paths to LUNs 5-2
- performance, RAID Group 2-14
- physical disk unit, *see* LUN (logical unit)
- physical volume, *see* LUN (logical unit)
- planning
 - hardware 6-10
 - LUNs and file systems 5-6
 - see also* worksheet
- plug types 6-10
- port, SP ALPA ID 5-17
- power requirements, storage system 6-9
- primary image
 - MirrorView 4-11
 - MirrorView/A 4-17

R

- rackmount storage system hardware components 6-3
- RAID 0 Group
 - defined 2-10
 - sample applications 2-21
 - when to use 2-19
- RAID 1 mirrored pair
 - defined 2-7
 - sample applications 2-20
 - when to use 2-18
- RAID 1/0 Group
 - defined 2-8
 - sample applications 2-21
 - when to use 2-18
- RAID 3 Group
 - defined 2-5
 - sample applications 2-20
 - when to use 2-18
- RAID 5 Group
 - defined 2-4
 - sample applications 2-20
 - when to use 2-18
- RAID Group
 - application guidelines 2-20
 - capacity 5-19
 - performance trade-offs 2-14

- planning worksheet 5-6
- sample applications 2-20
- shared storage examples 5-2
- type guidelines 2-18
- types, compared 2-13
- RAID Groups and LUNs 2-3
- read cache 6-6
- remote mirroring example
 - MirrorView 4-11
 - MirrorView/A 4-18
- replication over long distances 4-21
- revision testing, benefit with SnapView 4-3

S

- SAN Copy
 - benefit 4-27
 - benefits 4-27
 - example of 4-26
 - planning worksheets 4-30
 - simultaneous sessions 4-27
 - using with replication software 4-27
- server
 - cabling guidelines 6-10
 - connection to storage system 1-4
 - see also* cabling
- shared storage systems
 - application and LUN planning 5-7
 - components 6-8
 - examples 5-2
 - hardware planning worksheets 6-12
 - requirements 6-8
 - shared-or-clustered direct installation type, defined 1-9
- shared switched installation type
 - defined 1-9
 - sample 5-2
- site requirements, storage systems 6-9, 6-10
- size
 - chunk 4-9
 - RAID group 5-19
 - storage-system cache 5-17
 - storage-system components 6-9
- snapshot
 - compared to clone 4-7
 - overview 4-5

SnapView
 benefits 4-7
 integration with
 MirrorView 4-14
 MirrorView/A 4-20
 planning worksheet 4-8
 software mirroring 2-2
 SP (storage processor)
 about 6-5
 in remote mirror
 MirrorView 4-14
 IP address 5-17
 port ALPA ID 5-17
 storage-system cache size 5-17
 SPE (Storage Processor Enclosure)
 components 6-3
 dimensions and weight 6-9
 speed, disk 5-13
 storage
 system components 6-3
 types
 defined 1-9
 see also disk-array storage system, installation types
 storage system, *see* disk-array storage system
 stripe
 with RAID 1/0, RAID 0 2-8
 with RAID 5 2-4
 switch
 cascading 1-6
 description 1-6
 fabric, introduced 1-4
 introduced 1-4
 sample shared storage installation 5-2
 zone 1-6

T

trade-offs
 between snapshots and clones 4-7
 hardware, shared storage 6-8
 RAID types 2-13

U

unshared direct installation type, sample 5-5
 unshared storage systems
 application and LUN planning 5-7
 components 6-8
 defined 1-9
 hardware planning worksheets 6-12
 requirements 6-8
 sample 5-5

V

vault disks 2-11
 volume name 5-20

W

weight, storage system 6-9, 6-10
 worksheet
 hardware 6-12
 LUN 5-7, 5-14
 MirrorView 4-15
 MirrorView/A 4-22
 SAN Copy 4-30
 SnapView 4-8
 storage management 3-6
 write cache 6-6

Z

zone, switch 1-6

