

DAS – Disk Array Storage Systems

Planning a DAS Installation

Fibre Channel Environments

ORDER REFERENCE

86 A1 94JX 04



Bull DAS – Disk Array Storage Systems

Planning a DAS Installation
Fibre Channel Environments

Hardware

July 2000

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A1 94JX 04

The following copyright notice protects this book under the Copyright laws of the United States of America and other countries which prohibit such actions as, but not limited to, copying, distributing, modifying, and making derivative works.

Copyright © Bull S.A. 1992, 2000

Printed in France

Suggestions and criticisms concerning the form, content, and presentation of this book are invited. A form is provided at the end of this book for this purpose.

To order additional copies of this book or other Bull Technical Publications, you are invited to use the Ordering Form also provided at the end of this book.

Trademarks and Acknowledgements

We acknowledge the right of proprietors of trademarks mentioned in this book.

AIX[®] is a registered trademark of International Business Machines Corporation, and is being used under licence.

UNIX is a registered trademark in the United States of America and other countries licensed exclusively through the Open Group.

Year 2000

The product documented in this manual is Year 2000 Ready.

The information in this document is subject to change without notice. Groupe Bull will not be liable for errors contained herein, or for incidental or consequential damages in connection with the use of this material.

Preface

This planning guide provides an overview of CLARiiON® Fibre Channel disk-array storage system models and offers essential background information and worksheets to help you with the installation and configuration planning.

Please read this guide

- if you are considering purchase of a CLARiiON® Fibre Channel disk-array storage system and want to understand its features; or
- before you plan the installation of a storage system.

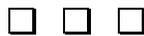
Audience for the manual

You should be familiar with the host servers that will use the storage systems and with the operating systems of 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 of the manual

- | | |
|-----------|--|
| Chapter 1 | Provides background information on the Fibre Channel protocols and explains the two major types of storage: shared and unshared. |
| Chapter 2 | Describes the RAID groups and the different ways they store data. |
| Chapter 3 | Describes the configurations for shared storage. |
| Chapter 4 | Describes the configurations for unshared storage. |
| Chapter 5 | Describes the hardware components for shared and unshared storage. |
| Chapter 6 | Describes storage-system management utilities. |



Contents

Chapter 1 – About Fibre Channel storage systems and networks (SANs)

Fibre Channel components	1-3
Server component (host interface kit with software)	1-3
Interconnect components	1-3
About shared storage and SANs (storage area networks)	1-7
Storage Groups	1-8
Storage system hardware for shared storage	1-10
About unshared storage	1-11
Storage system hardware for unshared storage	1-11
What next?	1-13

Chapter 2 – RAID types and tradeoffs

Introducing RAID	2-1
Disk striping	2-2
Mirroring	2-2
RAID groups and LUNs	2-2
RAID types	2-3
RAID 5 group (individual access array)	2-3
RAID 3 group (parallel access array)	2-4
RAID 1 mirrored pair	2-5
RAID 0 group (nonredundant array)	2-6
RAID 1/0 group (mirrored RAID 0 group)	2-6
Individual disk unit	2-7
Hot spare	2-8
RAID benefits and tradeoffs	2-9
Performance	2-10
Storage flexibility	2-11
Data availability and disk space usage	2-11
Guidelines for RAID groups	2-13
Sample applications for RAID types	2-14
What next?	2-16

Chapter 3 – Planning file systems and LUNs with shared storage in a SAN

Dual paths to LUNs	3-1
Sample shared storage configuration	3-1
Planning applications, LUNs, and Storage Groups	3-4

Application and LUN planning	3-5
Application and LUN planning worksheet	3-5
LUN and Storage Group planning worksheet	3-7
LUN Details Worksheet	3-9
LUN Details Worksheet	3-10
What next?	3-13

Chapter 4 – Planning LUNs and file systems with unshared storage

Dual SPs and paths to LUNs	4-1
Unshared storage system configurations	4-1
Single-server configurations	4-1
Dual-server configurations	4-2
Configurations with hubs	4-4
Sample unshared storage configurations	4-5
Single server example without hub	4-5
Dual-server example without hub	4-6
Multiple server example with hubs	4-7
Planning applications and LUNs - Unshared storage	4-8
Application and LUN planning	4-9
Application and LUN planning worksheet	4-9
LUN planning worksheet - Rackmount	4-11
LUN Details worksheet	4-13
LUN Details Worksheet	4-14
LUN Details Worksheet	4-15
What next?	4-19

Chapter 5 – Storage-system hardware

Hardware for shared storage	5-1
Storage hardware — rackmount DPE-based storage systems	5-2
Disks	5-2
Storage processor (SP)	5-3
Hardware for unshared storage	5-4
Types of storage system for unshared storage	5-4
Disks	5-6
Storage processor (SP)	5-7
Planning your hardware components	5-8
Configuration tradeoffs - shared storage	5-8
Configuration tradeoffs - unshared storage	5-9
Hardware data sheets	5-11
DPE data sheet	5-11
iDAE data sheet	5-13
DAE data sheet	5-14
30-slot system with SCSI disks data sheet	5-15

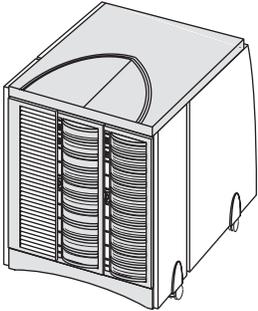
Cabinets for rackmount enclosures	5-16
Cable and configuration guidelines	5-16
Hardware planning worksheets	5-19
Hardware for shared storage	5-19
Hardware for unshared storage	5-22
What next?	5-27

Chapter 6 – Storage-system management

Managing shared or unshared storage systems using Navisphere Manager software	6-2
Managing unshared storage systems using Navisphere Supervisor software	6-3
Monitoring DAE-only storage systems (JBODs)	6-4
Storage management worksheets	6-4
Management Utility Worksheet – Shared Storage	6-4
Management Utility Worksheet – Unshared Storage	6-5

1

About Fibre Channel storage systems and networks (SANs)

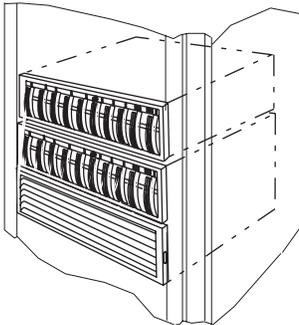


CLARiiON® Fibre Channel disk-array storage systems provide terabytes of disk storage capacity, high transfer rates, flexible configurations, and highly available data at low cost.

A storage system package includes a host interface kit with hardware and software to connect with a server, storage management software, Fibre Channel interconnect hardware, and one or more storage systems.

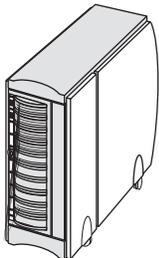
This chapter introduces Fibre Channel, its components, and shared and unshared storage systems. Major topics are

- Fibre Channel background
- Fibre Channel components
- About shared storage and SANs (storage area networks)
- About unshared storage



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 switched fabric point-to-point and arbitrated loop (FC-AL). The topologies used by the Fibre Channel storage systems described in this manual are switched fabric and FC-AL.

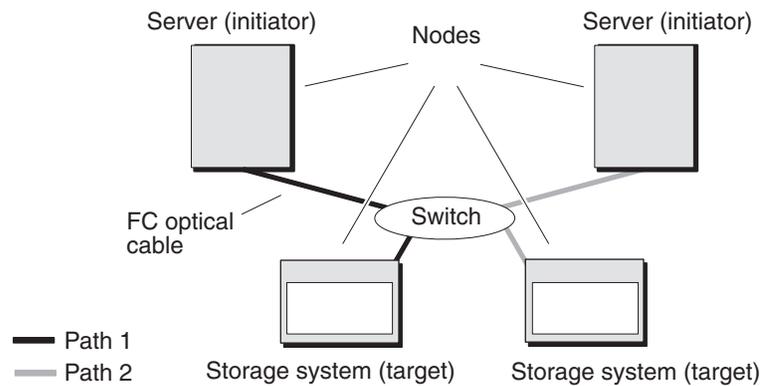


A switched fabric is set of point-to-point connections between nodes, the connection being made through a Fibre Channel switch. Each node may have its own unique address, but the path between nodes is governed by the switch. The nodes are connected by optical cable.

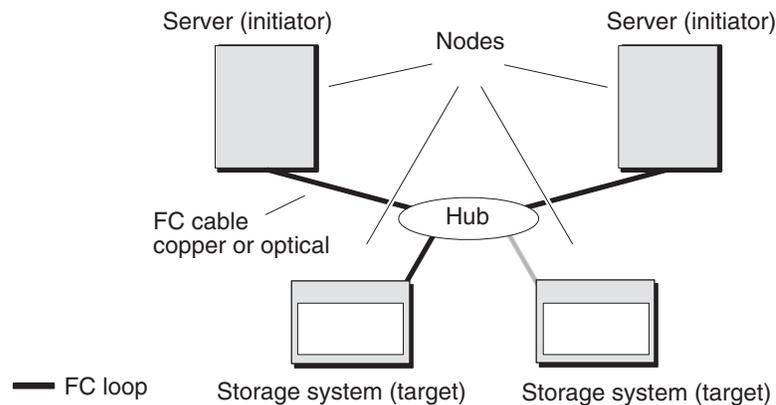
A Fibre Channel arbitrated loop is a circuit consisting of nodes. Each node has a unique address, called a Fibre Channel arbitrated loop address. The nodes are connected by copper or optical cables. An optical cable can transmit data over great distances for connections that span entire enterprises and can support remote disaster recovery systems. Copper cable serves well for local connections; its length is limited to 30meters (99 feet).

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

Nodes connected by switch - initiator and target



Nodes connected by hub - initiator and target



The Fibre Channel standards include rules about cable types and lengths.

Fibre Channel components

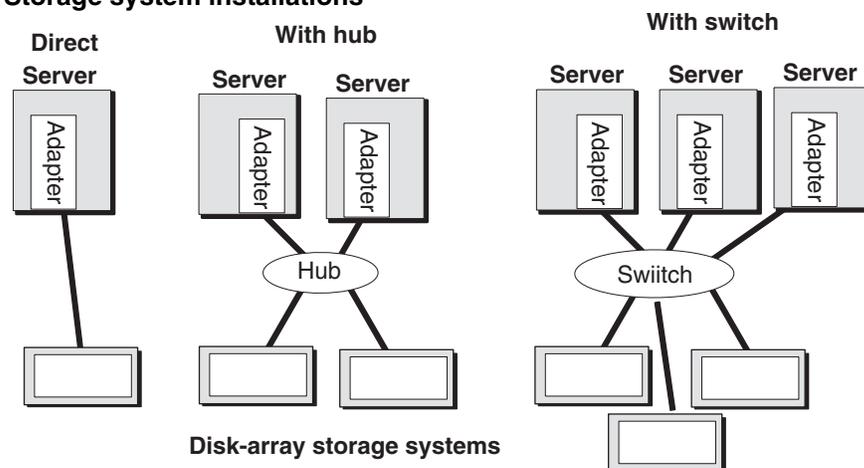
A Fibre Channel storage system has three main components:

- Server component (host interface kit with software)
- Interconnect components (cables and Fibre Channel switches and hubs)
- Storage components (storage systems and their hardware)

Server component (host interface kit with software)

The host interface kit includes a host-bus adapter (HBA) and support software. An HBA 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).

Storage system installations



Depending on your server type, you may have a choice of adapters. The adapter is designed for a specific host bus; for example, a PCI bus or SBUS. Some adapter types support copper or optical cabling; some support copper cabling only.

Interconnect components

The interconnect components include the cables, Fibre Channel switch (for shared storage), and Fibre Channel hub (for unshared storage).

Cables

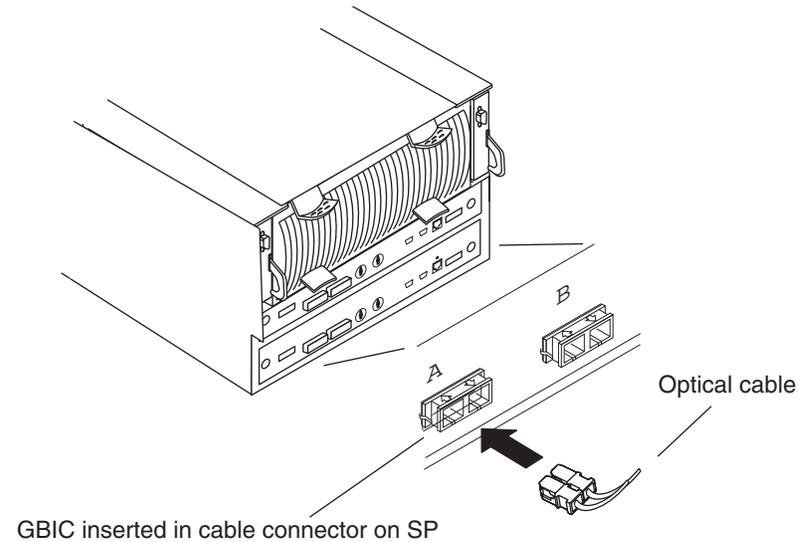
Depending on your needs, you can choose copper or optical cables.

The maximum length of copper cable is 30 meters (99 feet) between nodes or hubs. The maximum length of optical cable between server and hub or storage system is much greater, depending on the cable type. For example, 62.5 micron multimode cable can span up to 500 meters (1,640 feet) while 50

micron single-mode cable can span up to 10 kilometers (6.2 miles). This ability to span great distances is a major advantage of optical cable.

Some nodes have connections that require copper or optical cable. Other nodes allow for the conversion from copper to optical using a conversion device called a GigaBit Interface Converter (GBIC) or Media Interface Adapter (MIA). In most cases, a GBIC or MIA lets you substitute long-distance optical connections for shorter copper connections. A GBIC, inserted in the back of the storage system SP, looks like this:

GBIC (GigaBit Interface Converter) seen from back of storage system

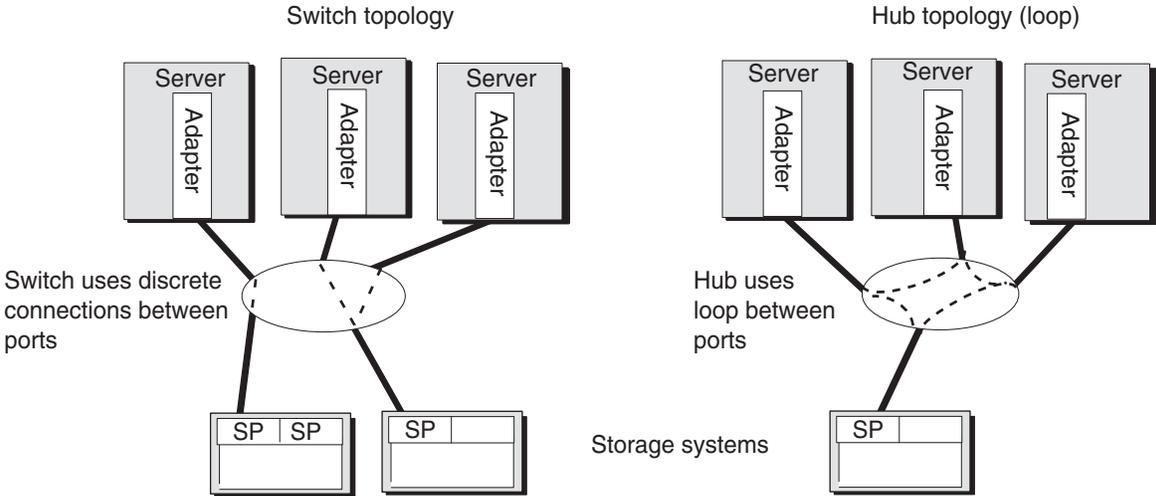


Details on cable lengths and rules appear later in this manual.

Fibre Channel switches

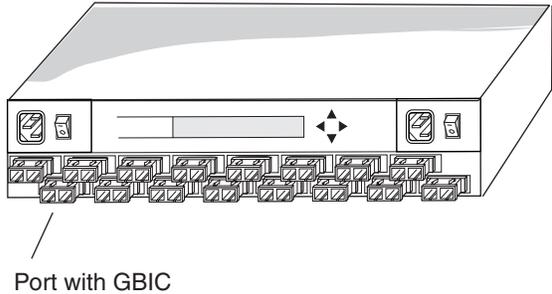
A Fibre Channel switch, which is a requirement for shared storage (a Storage Area Network, SAN) connects all the nodes cabled to it using a fabric topology. A switch adds serviceability and scalability to any installation; it allows on-line insertion and removal of any device on the fabric and maintains integrity if any connected device stops participating. A switch also provides host-to-storage-system access control in a multiple-host shared storage environment. A switch has several advantages over a hub: it provides point-to-point connections (as opposed to a hub's loop that includes all nodes) and it offers zoning to specify paths between nodes in the switch itself.

Switch and hub topologies compared



Fibre Channel switches are available with 16 or 8 ports. They are compact units that fit in 2 U (3.5 inches) for the 16-port or 1 U (1.75 inches) for the 8-port. They are available to fit into a rackmount cabinet or as small deskside enclosures.

16-port switch, back view

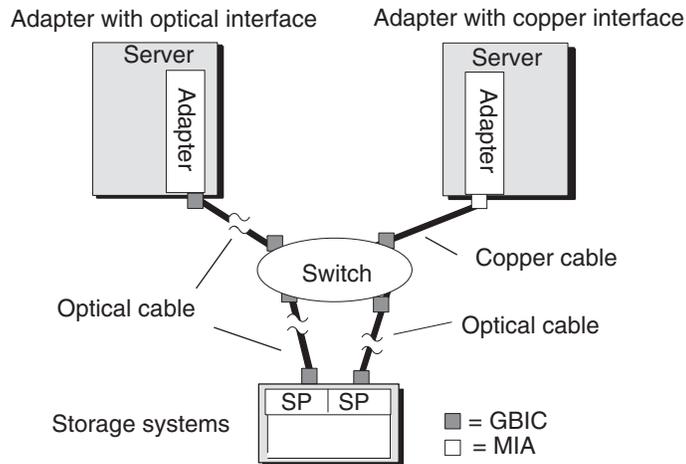


If your servers and storage systems will be far apart, you can place the switches closer to the servers or the storage systems, as convenient.

Switch interconnect options

Optical cable can span great distances between nodes. The SPs for shared storage support optical cable. And all shared storage systems ship with devices called GBICs (GigaBit Interface Converters) for use with optical cable. Many types of host-bus adapter support optical cable, and can thus use GBICs to connect to the cables. For adapters that require copper cable, a device called a MIA (Media Interface Adapter) lets the adapter communicate with the copper cable.

Adapters with optical and copper interfaces, shared storage



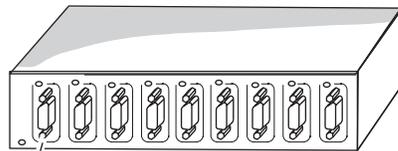
A switch is technically a repeater, not a node, in a Fibre Channel loop. However, it is bound by the same cabling distance rules as a node.

Fibre Channel hubs

A hub connects all the nodes cabled to it into a single logical loop. A hub adds serviceability and scalability to any loop; it allows on-line insertion and removal of any device on the loop and maintains loop integrity if any connected device stops participating.

Fibre channel hubs are compact units that fit in 1 U (1.75 inches) of storage space. They are available to fit into a rackmount cabinet or as small deskside units.

Nine-port hub



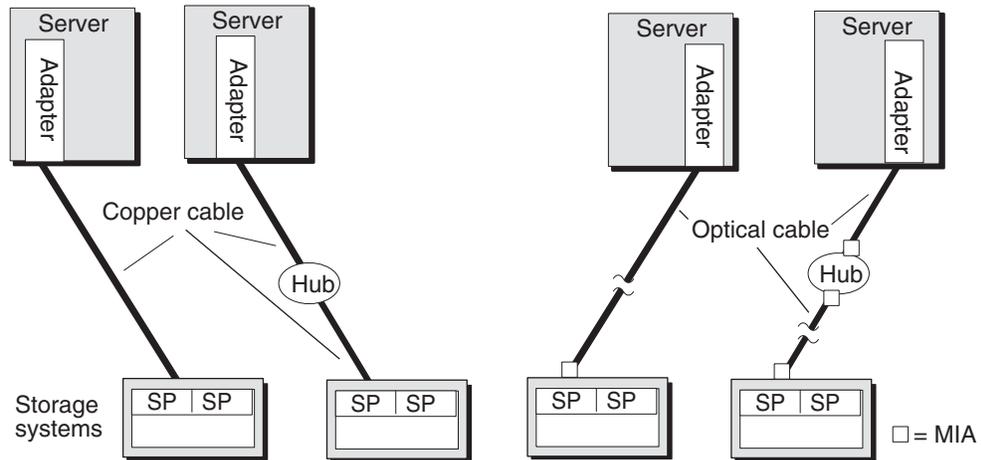
The nine-pin port can connect to a server, storage system, or another hub.

If your servers and storage systems will be far apart, you can place the hubs closer to the servers or the storage systems, as convenient. You can connect one hub to another, at great distances using optical cable, for disaster recovery systems.

Hub interconnect options

Some host-bus adapter types support copper cabling only; some support either optical or copper cabling. Hubs and disk-array storage systems can connect via copper cables or optical cables using MIAs.

Adapters with copper and optical interfaces, unshared storage



A hub 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.

About shared storage and SANs (storage area networks)

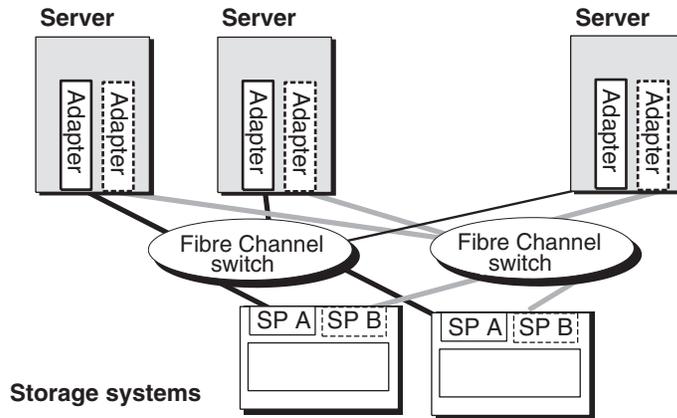
This section explains the features that let multiple servers share disk-array storage systems on a SAN (storage area network).

A SAN is a collection of storage devices connected to servers via 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

A SAN includes two or more servers and has Fibre Channel switches between servers and storage systems.

Components of a SAN



Fibre Channel switches can control data access to storage systems through the use of switch zoning. With zoning, an administrator can specify groups (called zones) of Fibre Channel devices (such as host-bus adapters, specified by worldwide name), and SPs between which the switch will allow communication.

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. 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.

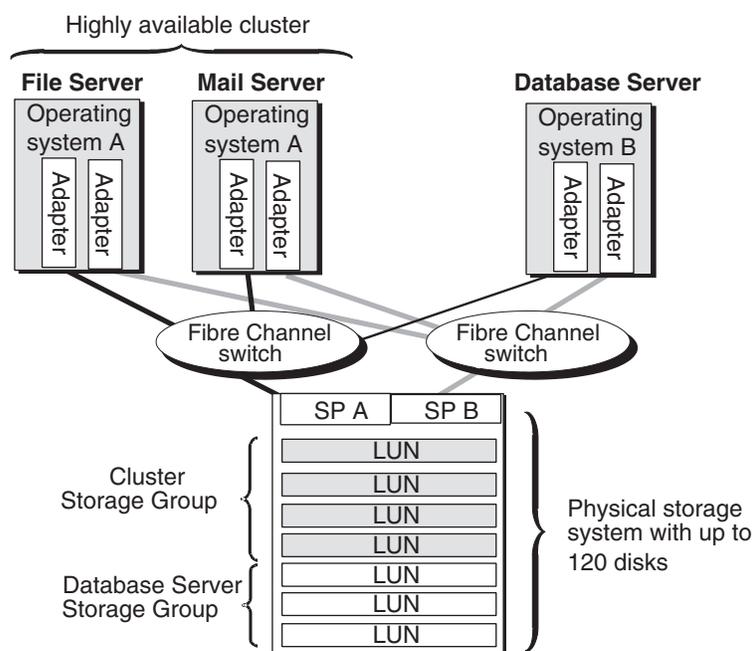
Storage Groups

A Storage Group is one or more LUNs (logical units) within a disk-array storage system that is reserved for one or more servers and is inaccessible to other servers.

When you configure the storage system for a SAN, you specify servers and the Storage Group(s) each server can read from and/or write to. The Licensed Internal Code firmware running in each storage system enforces the server-to-Storage Group permissions.

The following figure shows a simple SAN configuration consisting of one CLARiiON 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.

Sample SAN configuration



Access Control in a SAN

Access control permits or restricts a server's access to SAN storage. There are two kinds of access control:

- Data access control
- Configuration access control

Data access control — Storage Groups provide data access control. During storage system configuration, using a management utility, the system administrator associates a server with one or more LUNs. The associated LUNs constitute a Storage Group. The server HBA port's unique worldwide name identifies the port and allows access to the Storage Group. The storage-system SP will prevent any server from accessing LUNs the server is not configured to access.

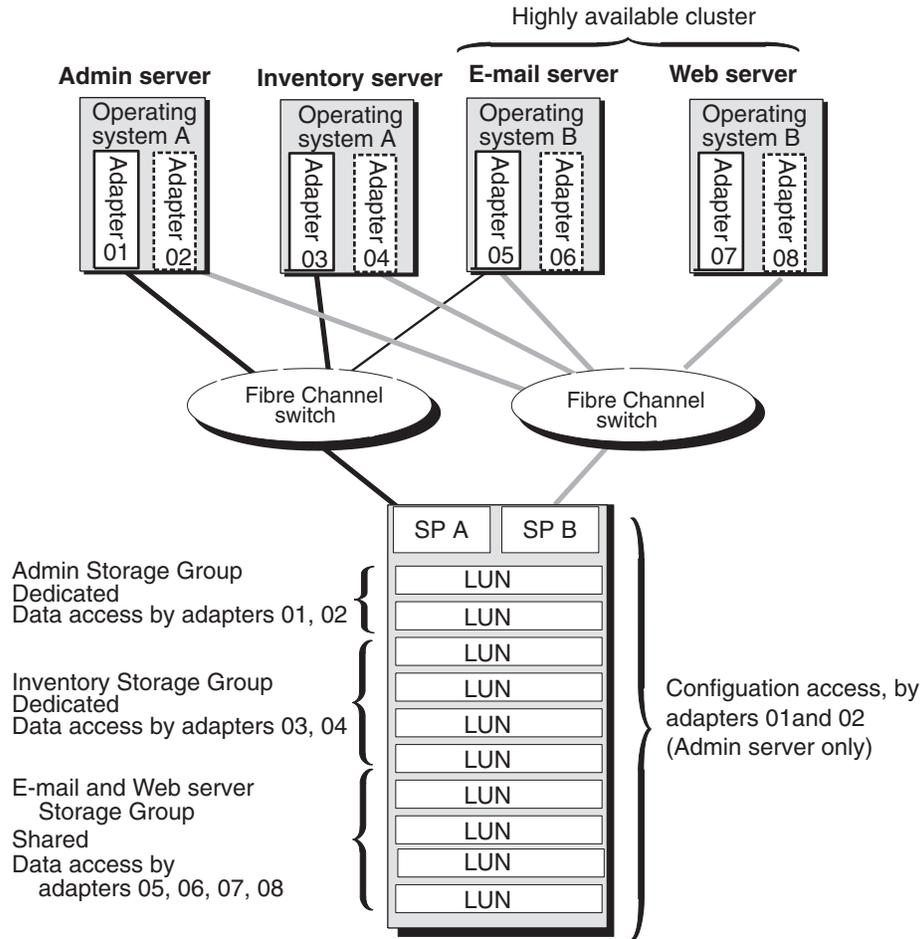
Each server sees its Storage Group as if it were an entire storage system and never sees the other LUNs on the storage system. Therefore, it cannot access or modify data on LUNs that are not part of its Storage Group.

Configuration access control — Configuration access control lets you restrict the servers through which a user can send configuration commands to an attached storage system. Configuration access is governed by a password. The administrator can set the password during storage system setup; or if the site doesn't want to use configuration access control, it can skip the password.

The following figure shows both data access control (Storage Groups) and configuration access control. Each server has exclusive read and write access to its designated Storage Group. Of the four servers connected to the

SAN, only the Admin server can send configuration commands to the storage system.

Data and configuration access control in a SAN

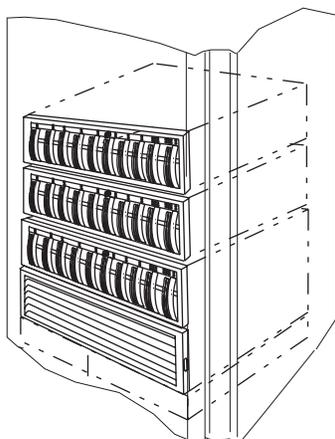


Storage system hardware for shared storage

For shared storage, you need a Disk-array Processor Enclosure (DPE) storage system.

A DPE is a 10-slot enclosure with hardware RAID features provided by one or two storage processors (SPs). For shared storage, two SPs are required. In addition to its own disks, a DPE can support up to 11 ten-slot Disk Array Enclosures (DAEs) for a total of 120 disks.

Disk-array processor enclosure (DPE) with two DAEs, for shared storage



About unshared storage

Unshared storage systems are less costly and less complex than shared storage systems. They offer many shared storage system features; for example, you can use multiple unshared storage systems with multiple servers. However, with multiple servers, unshared storage offers less flexibility and security than shared storage, since any user with write access to privileged server files can enable access to any storage system.

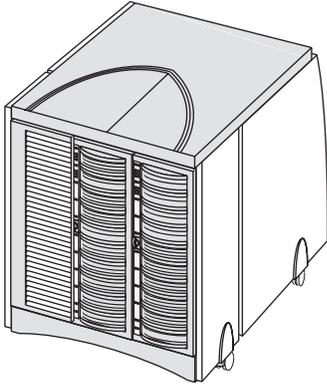
Storage system hardware for unshared storage

For unshared storage, there are four types of storage system, each using the FC-AL protocol. Each type is available in a rackmount or desktside (office) version.

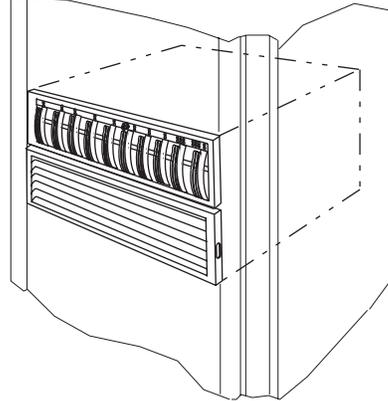
- **Disk-array Processor Enclosure (DPE) storage systems.** A DPE is a 10-slot enclosure with hardware RAID features provided by one or two storage processors (SPs). In addition to its own disks, a DPE can support up to 110 additional disks in 10-slot Disk Array Enclosures (DAEs) for a total of 120 disks. This is the same type of storage system used for shared storage, but it has a different SP and different Licensed Internal Code (LIC).
- **Intelligent Disk Array Enclosure (iDAE).** An iDAE, like a DPE, has SPs and thus all the features of a DPE, but is thinner and has a limit of 30 disks.
- **Disk Array Enclosure (DAE).** A DAE does not have SPs. A DAE can connect to a DPE or an iDAE, or you can use it without SPs. A DAE used without an SP does not inherently include RAID, but can operate as a RAID device using software running on the server system. Such a DAE is also known as Just a Box of Disks, or JBOD.
- **30-slot SCSI-disk storage systems.** Like the DPE, these offer RAID features provided by one or two SPs. However, they use SCSI, not Fibre Channel, disks. Each has space for 30 disks.

Disk-array processor enclosure (DPE)

Deskside DPE with DAE

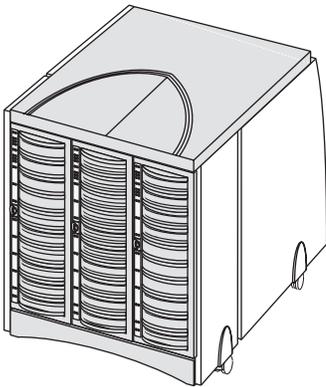


Rackmount DPE, one enclosure, supports up to 11 DAEs

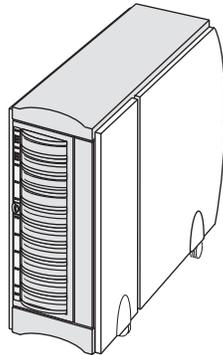


Intelligent disk-array enclosure (IDAE)

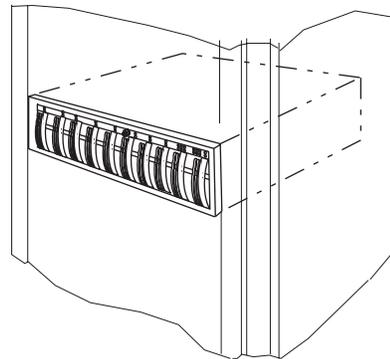
30-slot deskside



10-slot deskside

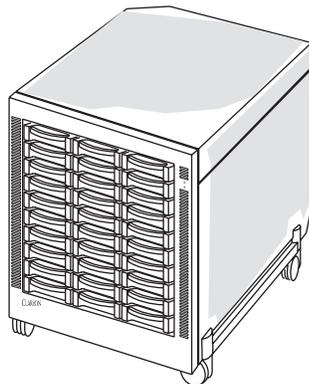


Rackmount

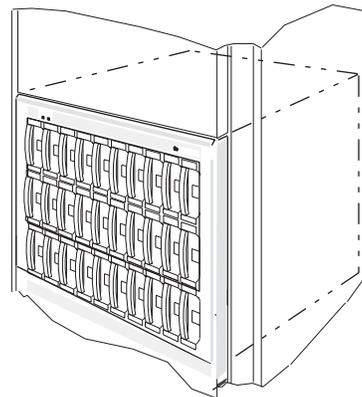


30-slot SCSI-disk storage system

Deskside



Rackmount



What next?

For information about RAID types and RAID tradeoffs, continue to the next chapter. To plan LUNs and file systems for shared storage, skip to Chapter 3; or for unshared storage, Chapter 4. For details on the storage-system hardware -- shared and unshared -- skip to Chapter 5.

For storage-system management utilities, skip to Chapter 6.



2

RAID types and tradeoffs

This chapter explains 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 skip to the planning chapter (Chapter 3 for shared storage; Chapter 4 for unshared storage). Topics are as follows:

- Introducing RAID
- RAID 5 group (individual access array)
- RAID 3 group (parallel access array)
- RAID 1 mirrored pair
- RAID 0 group (nonredundant array)
- RAID 1/0 group (mirrored RAID 0 array)
- Individual disk
- Hot spare
- RAID benefits and tradeoffs
- Guidelines for RAID groups
- Sample applications for RAID groups

IMPORTANT This chapter applies primarily to storage systems with storage processors (SPs). For a storage system without SPs (a DAE-only system), RAID types are limited by the RAID software you run on the server. The RAID terms and definitions used here conform to generally accepted standards.

Introducing RAID

The storage system uses RAID (redundant array of independent disks) technology. RAID technology groups separate disks into one logical storage 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 groups 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) and a five-disk group. The group has five disks, so 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

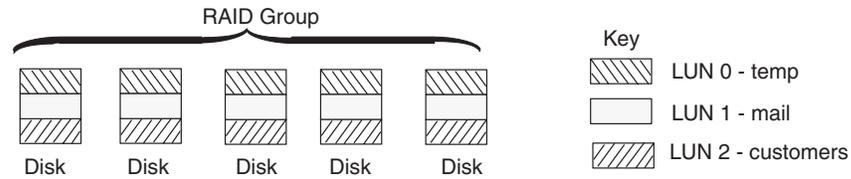
Mirroring maintains a second (and optionally through software, a third) copy of a logical disk *image* that provides continuous access if the original image becomes inaccessible. The system and user applications continue running on the good image without interruption. There are two kinds of mirroring: hardware mirroring, in which the SP synchronizes the disk images; and software mirroring, in which the operating system synchronizes the images.

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. Or you can use software mirroring with RAID 0 groups or individual units that have no inherent data redundancy. With software mirroring, the operating system mirrors the images. Some operating systems support software mirroring; others do not.

RAID groups and LUNs

On full-fibre systems, 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 36-Gbyte disks offers 144 Gbytes of space. You could bind three LUNs, say with 24, 60, and 60 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 degrade. 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.



RAID types

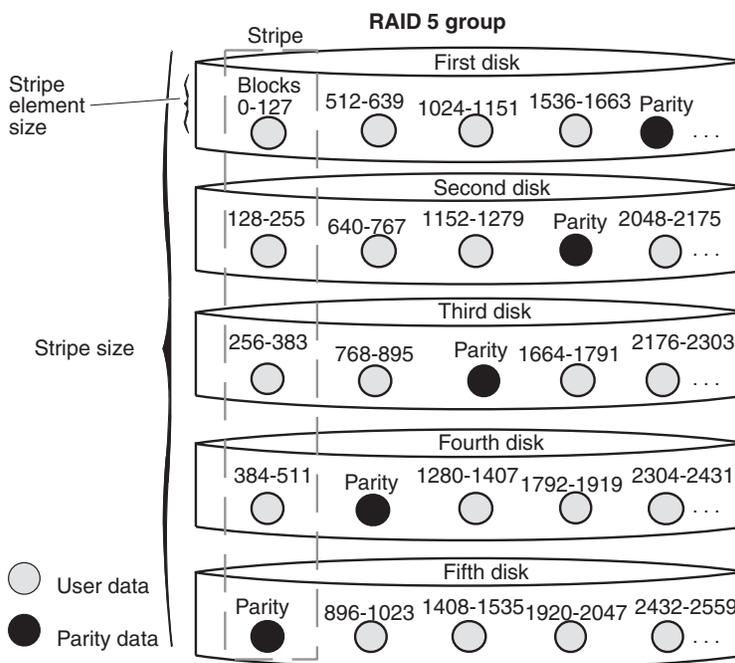
You can choose from the following RAID types: RAID 5, RAID 3, RAID 1, RAID 0, RAID 1/0, individual disk unit, and hot spare.

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 on a full-fibre storage system, you can create up to 32 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.

The following figure 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.



The storage system performs more steps when writing data to a RAID 5 group than to any other type of group. For each write, the storage system

1. Reads user data from the sectors and parity data for the sectors.
2. Recalculates the parity data.
3. Writes the new user and parity data.

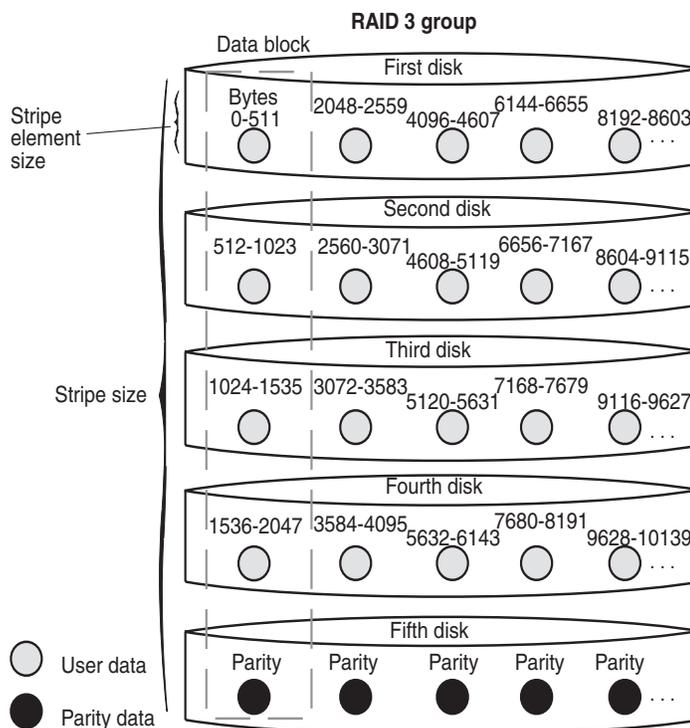
RAID 5 groups benefit greatly from storage-system caching, particularly write caching.

RAID 3 group (parallel access array)

A RAID 3 group consists of five or more disk disks. The hardware always reads from or writes to all the disks. A RAID 3 group uses disk striping. To maintain the RAID 3 performance, you can create only one LUN per RAID 3 group.

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.

The following figure 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.



RAID 3 differs from RAID 5 in several important ways. First, in a RAID 3 group the hardware processes disk requests serially; whereas in a RAID 5 group the hardware can interleave disk requests. Second, with a RAID 3 group, the parity information is stored on one disk; with a RAID 5 group, it is stored on all disks. Finally, with a RAID 3 group, the I/O occurs in small units (one sector) to each disk. A RAID 3 group works well for single-task applications that use I/Os of blocks larger than 64 Kbytes.

Each RAID 3 group requires some dedicated SP memory (6 Mbytes recommended per group). This memory is allocated when you create the group and becomes unavailable for storage-system caching. For top performance, we suggest that you do not use RAID 3 groups with RAID 5, RAID 1/0, or RAID 0 groups, since SP resources (including memory) are best devoted to the RAID 3 groups. RAID 1 mirrored pairs and individual units require less SP attention and therefore work well with RAID 3 groups.

For each write to a RAID 3 group, the storage system

1. Calculates the parity data.
2. Writes the new user and parity data.

RAID 1 mirrored pair

A RAID 1 group consists of two disks that are mirrored automatically by the storage-system hardware.

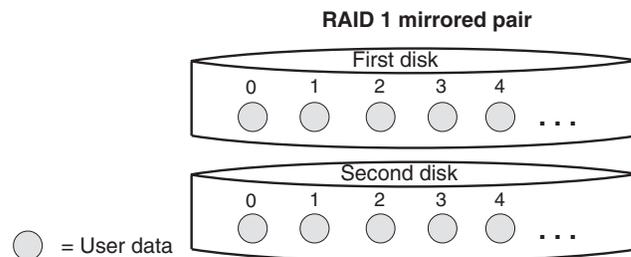
RAID 1 hardware mirroring within the storage system is not the same as software mirroring or hardware mirroring for other kinds of disks.

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 in 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.



RAID 0 group (nonredundant array)

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. In a full-fibre storage system, you can create up to 32 LUNs per RAID 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. RAID 0 offers enhanced performance through simultaneous I/O to different disks.

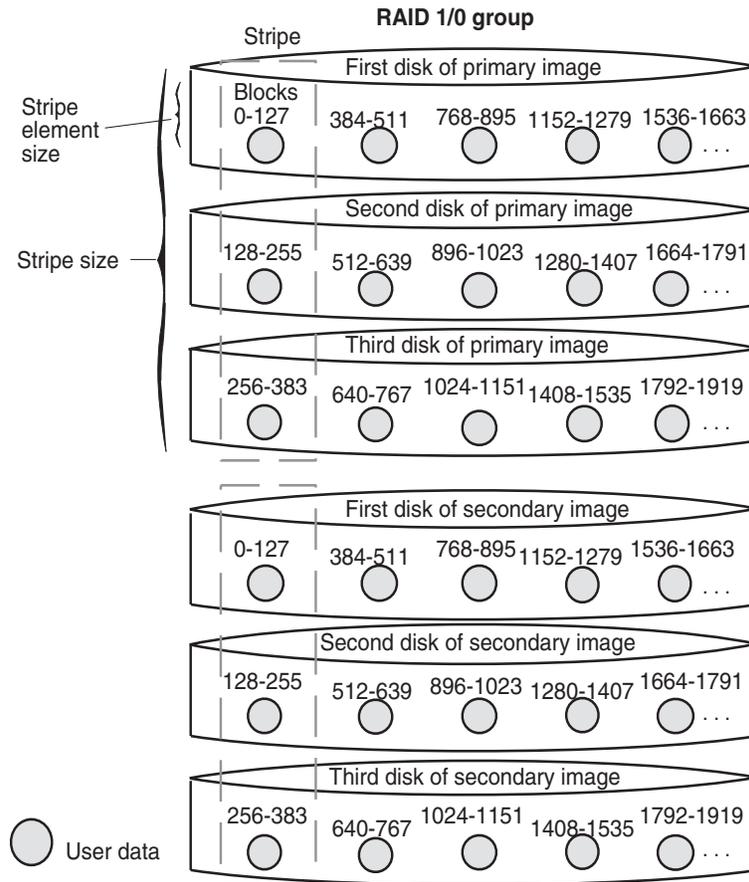
If the operating system supports software mirroring, you can use software mirroring with the RAID 0 group to provide high availability. A desirable alternative to RAID 0 is RAID 1/0.

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 five 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 on a

full-fibre storage system, you can create up to 32 RAID 5 LUNs to apportion disk space to different users, servers, and applications.

The following figure 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.



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

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.

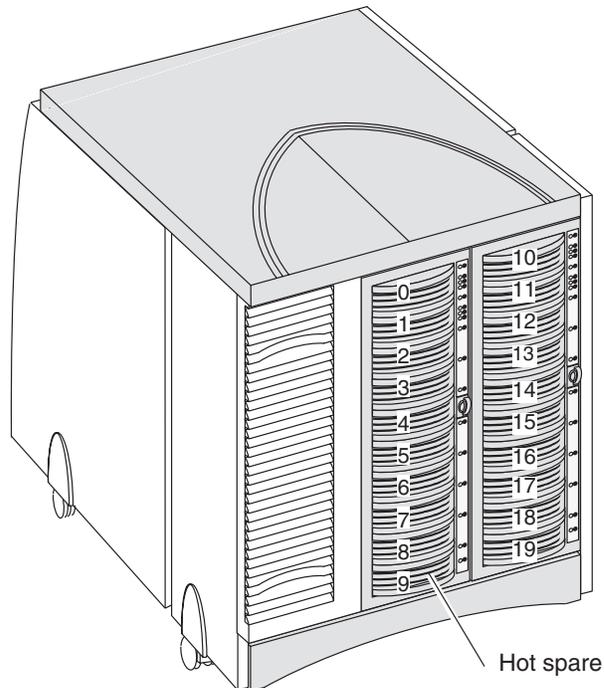
When the copy is done, the disk group consists of disks in the original slots, and the SP automatically frees the hot spare to serve as a hot spare again. A hot spare is most useful when you need the highest data availability. It eliminates the time and effort needed for someone to notice that a disk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

IMPORTANT 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 for a disk that serves for Licensed Internal Code storage or the write cache vault. That is, a hot spare can be any of the following disks:

DPE or iDAE system without write caching:	disks 3-119
DPE system with write caching:	disks 9-119
iDAE system with write caching:	disks 5-29
30-slot SCSI-disk system:	disks A1-E1, A2-E2, B3-E3, A4-E4

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



1. RAID 5 group consists of disk modules 0-4; RAID 1 mirrored pair is modules 5 and 6; hot spare is module 9.
2. Disk module 3 fails.
3. RAID 5 group becomes modules 0, 1, 2, 9, and 4; now no hot spare is available.
4. System operator replaces failed module 3 with a functional module.
5. RAID 5 group once again is 0-4 and hot spare is 9.

RAID benefits and tradeoffs

This section reviews RAID types and explains their benefits and tradeoffs. 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, but can be software mirrored if the operating system supports mirroring
- Individual unit; no inherent high-availability features but can be software mirrored, if the operating system supports mirroring
- 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

IMPORTANT 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.

The following table compares the read and write performance, tolerance for disk failure, and relative cost per megabyte (Mbyte) of the RAID types and their software-mirrored counterparts. Figures shown are theoretical maximums.

Relative performance, availability, and cost of disk RAID types (individual unit = 1.0)

Disk configuration	Relative read performance without cache	Relative write performance without cache	Relative cost per Mbyte
RAID 5 group with five disks	Up to 5 with five disks (for small I/O requests, 2 to 8 Kbytes)	Up to 1.25 with five disks (for small I/O requests, 2 to 8 Kbytes)	1.25
RAID 3 group with five disks	Up to 4 (for large I/O requests)	Up to 4 (for large I/O requests)	1.25
RAID 1 mirrored pair	Up to 2	Up to 1	2
RAID 1/0 group with 10 disks	Up to 10	Up to 5	
Individual unit			
Not software mirrored	1	1	1
Software mirrored	Up to 2	Up to 1	2
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 load on server and storage system.			

Performance

RAID 5, with individual access, provides high read throughput for small requests (blocks of 2 to 8 Kbytes) by allowing simultaneous reads from each disk in the group. RAID 5 write throughput is limited by the need to perform four I/Os per request (I/Os to read and write data and parity information). However, write caching improves RAID 5 write performance.

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 five disks in each request but need not read data and parity before writing – advantageous for large requests but not for small ones. RAID 3 employs SP memory without caching, which means you do not need the second SP and BBU 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 1Mbyte. 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 six disks; a RAID 0 group, at least three disks.

An individual unit needs 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 32 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.

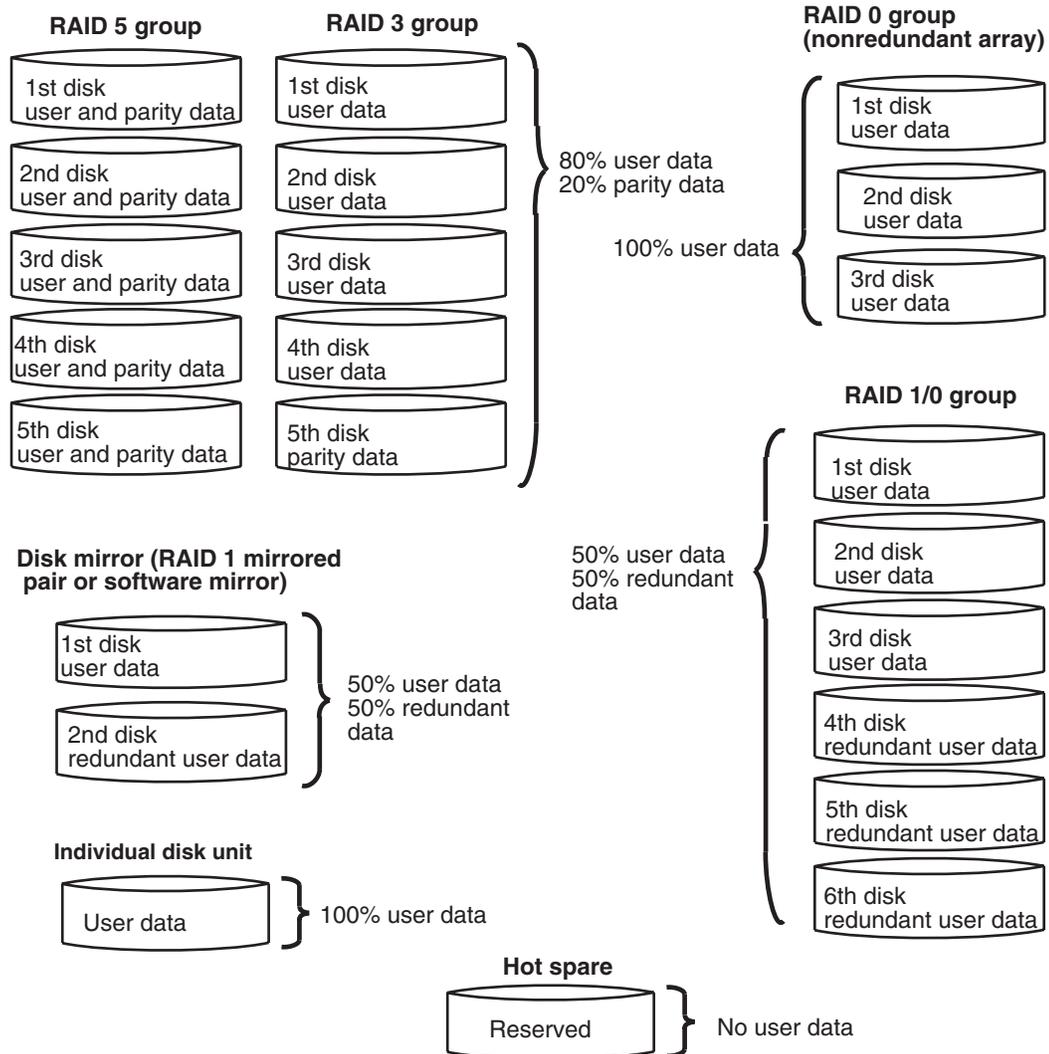
Data availability and disk space usage

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. Or bind a RAID 0 group or individual disk unit that you will later mirror with software mirroring. If data availability is not critical, or disk space usage is critical, bind an individual unit or RAID 0 group without software mirroring.

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 on page 2-10.

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 of five disks, 80 percent of the disk space is available for user data. 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.

Disk space usage in the RAID configurations



If your operating system supports software mirroring, you can use software mirroring for an individual unit or RAID 0 group. For a comparison of software and hardware mirroring, see the section on mirroring, page 2-2.

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 or use software mirroring with the RAID 0 group. A RAID 1/0 group or software-mirrored RAID 0

group provides the best combination of performance and availability, at the highest cost per Mbyte of disk space. A RAID 1/0 group offers higher reliability than a software-mirrored RAID 0 group.

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 the figure on page 2-12. For high availability with an individual disk, you can use software mirroring.

Guidelines for RAID groups

To decide when to use a RAID 5 group, RAID 3 group, mirror (that is, a RAID 1 mirrored pair, RAID 1/0 group, or software mirroring), 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
- Good read and moderate write performance are important (write caching can improve RAID 5 write performance)
- 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 (or you plan to use software mirroring by the operating system)
- Overall performance is very important

Use an individual unit for applications where

- High availability is not important (or you plan to use software mirroring by the operating system)
- 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 multitask 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 or individual unit to be software mirrored — A RAID 1 mirrored pair or individual unit to be software mirrored 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. In terms of high availability, a RAID 0 group is less available than an individual unit. You can improve availability by software mirroring the RAID 0 group if the operating system supports mirroring. A RAID 0 group (like a RAID 5 group) requires a minimum of three disks; software mirroring a RAID 0 group requires the same number of disks used in the original image. A RAID 0 group is useful for applications using short-term data to which you need quick access, such as mail.

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 six 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. Individual disk units are flexible; later, if the operating system supports software mirroring, you can always use software mirroring with the disk.

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.

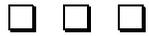
Hot spare — A hot spare provides no data storage but enhances the availability of each RAID 5, RAID 3, RAID 1, and RAID 1/0 group in a storage system. Use a hot spare where you must regain high availability quickly without human intervention if any disk in such a RAID group fails. A hot spare also minimizes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

What next?

What next?

This chapter explained RAID group types and tradeoffs. To plan LUNs and file systems for shared storage, continue to Chapter 3; or for unshared storage, skip to Chapter 4. For details on the storage-system hardware — shared and unshared — skip to Chapter 5.

For storage-system management utilities, skip to Chapter 6.



3

Planning file systems and LUNs with shared storage in a SAN

This chapter shows a sample RAID, LUN, and Storage Group configuration with shared storage, and then provides worksheets for planning your own shared storage installation. Topics are

- Paths to LUNs
- Sample shared storage configuration
- Shared storage planning worksheets

Dual paths to LUNs

A shared storage system includes one or more servers, two Fibre Channel switches, one or more storage systems, each with two SPs, and the shared model of storage system software (Licensed Internal Code).

With shared storage, there are two paths to each LUN in the storage system. The shared storage LIC detects both paths and, using a built-in application called Application Transparent Failover (ATF), can automatically switch to the other path if a device (like a host-bus adapter or cable) fails.

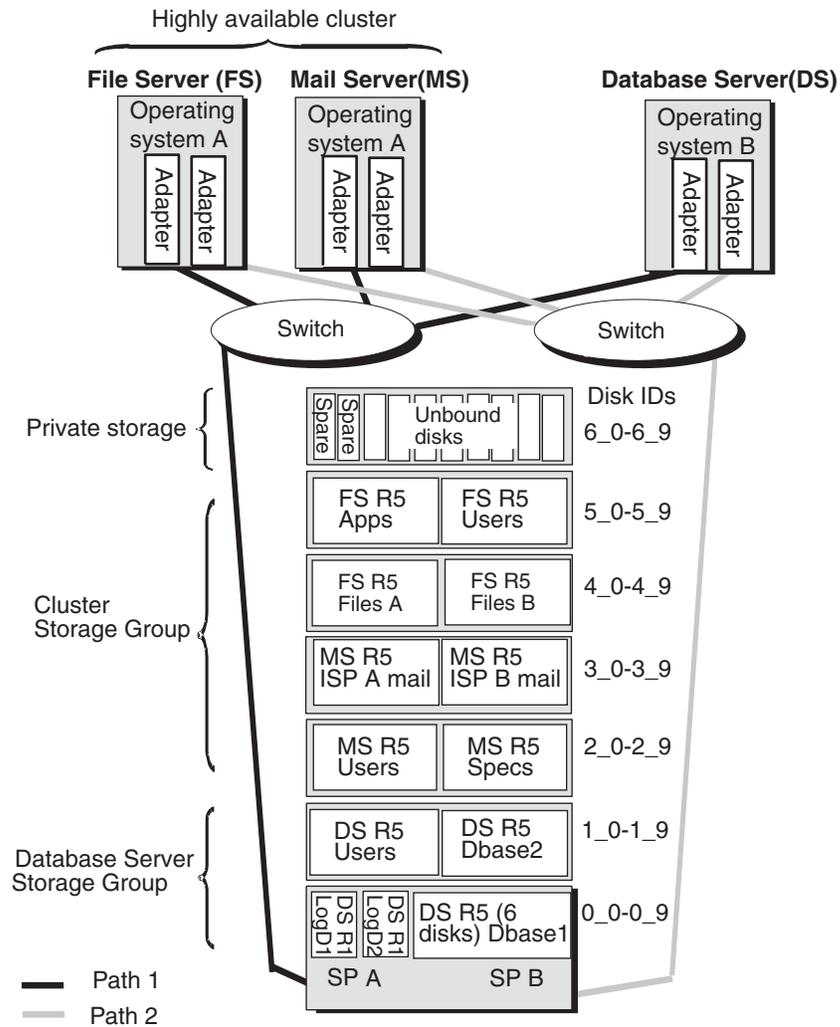
The SP on which you bind a disk becomes the default owner of the resulting LUN. The route through the SP that owns a unit determines the primary route to that LUN. The route through the other SP is the secondary route to the LUN.

With a single server, you can bind disks to each SP to share the I/O load equally. With dual servers, each server has one primary SP; the SP that binds a LUN determines the server that's the primary owner of the LUN.

Sample shared storage configuration

The following figure shows a sample shared storage system connected to three servers: two servers in a cluster and one server running a database management program.

Sample shared storage configuration



The storage system disk IDs and servers' Storage Group LUNs are as follows.

Clustered System LUNs		Database Server LUNs (DS) - SP A
File Server LUNs (FS) - SP B	Mail Server LUNs (MS) - SP A	
<u>Disks RAID type, storage type</u>	<u>Disks RAID type, storage type</u>	<u>Disks RAID type, storage type</u>
5_0-5_4 – RAID 5, Applications	2_0-2_4 – RAID 5, ISP A mail	0_0,0_1 – RAID 1, Log file for database Dbase1
5_5-5_9 – RAID 5, Users	2_5-2_9 – RAID 5, ISP B mail	0_2, 0_3 – RAID 1, Log file for database Dbase2
4_0-4_4 – RAID 5, Files A	3_0-3_4 – RAID 5, Users	0_4-0_9 – RAID 5 (6 disks), Dbase1
4_5-4_9 – RAID 5, Files B	3_5-3_9 – RAID 5, Specs	1_0-1_4 – RAID 5, Users
		1_5-1_9 – RAID 5, Dbase 2
6_0, 6_1 – Hot spare (usable for any server's LUN)		

With 18-Megabyte disks, the LUN storage capacities and drive names are as follows.

File Server — 288 Gbytes on four LUNs

FS R5 Apps Unit **S** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for applications.

FS R5 Users Unit **T** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for user directories and files.

FS R5 FilesA Unit **U** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for file storage.

FS R5 FilesB Unit **V** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for file storage.

Mail Server — 288 Gbytes on four LUNs

MS R5 ISP mail Unit **O** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for the mail delivered via ISP A.

MS R5 ISP mail Unit **P** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for the mail delivered via ISP B.

MSR5 Users Unit **Q** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for user directories and files.

MS R5 Specs Unit **R** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for specifications.

Database Server — 208 Gbytes on four LUNs

DS R5 Users Unit **users** on five disks bound as a RAID 5 group for 72 Gbytes of storage; for user directories.

DS R5 Dbase2 Unit **dbase2** on five disks bound as a RAID 5 group for 52 Gbytes of storage; for the second database system.

DS R1 Logs Unit **logfiles** on two disks bound as a RAID 1 mirrored pair for 18 Gbytes of storage; for the database log files.

DS R5 Dbase1 Unit **dbase** on eight disks bound as a RAID 5 group for 144 Gbytes of storage; for the primary database system.

Planning applications, LUNs, and Storage Groups

This section helps you plan your shared storage use — the applications to run, the LUNs that will hold them, and the storage groups that will belong to each server. The worksheets to help you do this 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 shared storage system.

Sample worksheets appear later in this chapter.

software to assign a volume label that describes the RAID configuration. 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 partition, file system, or logical volume. 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, then add a factor for growth.

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 following worksheet to select the disks that will make up the LUNs and Storage Groups in the SAN. A shared storage system can include up to 120 disks, numbered 0 through 119, left to right from the bottom up.

LUN and Storage Group planning worksheet

11_0	11_1	11_2	11_3	11_4	11_5	11_6	11_7	11_8	11_9
------	------	------	------	------	------	------	------	------	------

10_0	10_1	10_2	10_3	10_4	10_5	10_6	10_7	10_8	10_9
------	------	------	------	------	------	------	------	------	------

9_0	9_1	9_2	9_3	9_4	9_5	9_6	9_7	9_8	9_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

8_0	8_1	8_2	8_3	8_4	8_5	8_6	8_7	8_8	8_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

7_0	7_1	7_2	7_3	7_4	7_5	7_6	7_7	7_8	7_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

6_0	6_1	6_2	6_3	6_4	6_5	6_6	6_7	6_8	6_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

5_0	5_1	5_2	5_3	5_4	5_5	5_6	5_7	5_8	5_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

4_0	4_1	4_2	4_3	4_4	4_5	4_6	4_7	4_8	4_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

3_0	3_1	3_2	3_3	3_4	3_5	3_6	3_7	3_8	3_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

2_0	2_1	2_2	2_3	2_4	2_5	2_6	2_7	2_8	2_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

1_0	1_1	1_2	1_3	1_4	1_5	1_6	1_7	1_8	1_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

0_0	0_1	0_2	0_3	0_4	0_5	0_6	0_7	0_8	0_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

Storage system number or name: _____

Storage Group ID or name: _____ Server hostname: _____ Dedicated Shared

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: _____ Dedicated Shared

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: _____ Dedicated Shared

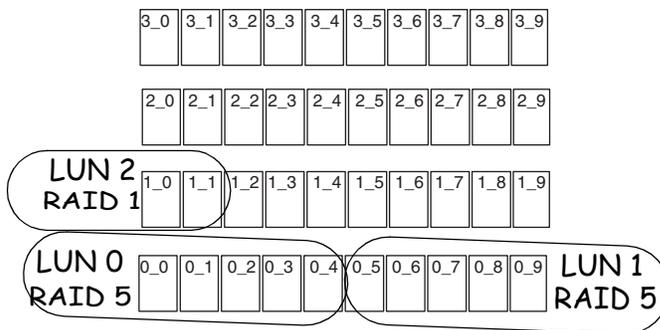
LUN ID or name _____ RAID type ___ Cap. (Gb) _____ Disk IDs _____

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 follows.



Storage system number or name: SS1

Storage Group ID or name: Mail 1 Server hostname: Server1 Dedicated Shared

LUN ID or name 0 RAID type 5 Cap. (Gb) 72 Disk IDs 0 0, 0 1, 0 2, 0 3, 0 4

LUN ID or name 1 RAID type 5 Cap. (Gb) 72 Disk IDs 0 5, 0 6, 0 7, 0 8, 0 9

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: Index1 Server hostname: Server2 Dedicated Shared

LUN ID or name 2 RAID type 18 Cap. (Gb) _____ Disk IDs 1 0, 1 1

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 _____

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. This form is *enclosure_diskID*, where *enclosure* is the enclosure number (the bottom one is 0, above it 1, and so on) and *diskID* is the disk position (left is 0, next is 1, and so on).

IMPORTANT None of the disks 0_0 through 0_8 may be used as a hot spare.

Next, complete as many of the Storage System sections as needed for all the Storage Groups in the SAN. Copy the (blank) worksheet as needed for all Storage Groups in each storage system.

A storage system is any group of enclosures connected to a DPE; it can include up to 11 DAE enclosures for a total of 120 disks. If a Storage Group will be dedicated (not accessible by another system 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.

LUN Details worksheet

Use the following LUN details worksheet to plan the individual LUNs. Complete as many of these as needed for all LUNs in your SAN.

Completing the LUN details worksheet

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

Storage-system entries

Storage-system configuration. specify Shared storage.

SP FC-AL address ID. This does not apply to shared storage, in which the switch determines the address of each device.

Use memory for caching. You can use SP memory for read/write caching or RAID 3. (Using both caching and RAID 3 in the same storage system is not recommended.) If you choose caching, check the box and continue to the next step; for RAID 3, skip to the RAID Group ID entry.

Read cache size. If you want a read cache, it should generally be about one third of the total available cache memory.

Write cache size. The write cache should be two thirds of the total available. Some memory is required for system overhead, so you cannot determine a precise figure at this time. For example, for 256 Mbytes of total memory, you might have 240 Mbytes available, and you would specify 80 Mbytes for the read cache and 160 Mbytes for the write cache.

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

- For a general-purpose file server — 8 Kbytes
- For a database application — 2 or 4 Kbytes

The ideal cache page size depends on the operating system and application.

Use memory for 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.

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:

RAID5 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 18-Gbyte disks holds 72 Gbytes;
- An eight-disk RAID 1/0 group of 18-Gbyte disks also holds 72Gbytes;
- A RAID 1 mirrored pair of 18-Gbyte disks holds 18 Gbytes; and
- An individual disk of an 18-Gbyte disk also holds 18 Gbytes.

Each disk in the RAID group must have the same capacity; otherwise, you will waste disk storage space.

LUN size. Enter the user-available capacity in gigabytes (Gbytes) of the LUN. You can make this the same size as the RAID group, above. 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 ID(s) 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 the DPE (enclosure 0 disks 2 through 6, enter 0_2, 0_3, 0_4, 0_5, and 0_6.

SP. 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.

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 3-11), you can specify whether you want caching — 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.

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. Enter the name of each server (copied from the LUN and Storage Group worksheet).

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 this storage system’s chassis will be shared with another server, and the other server is the primary owner of this disk, 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?

This chapter outlined the planning tasks for shared storage systems. 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 5.



4

Planning LUNs and file systems with unshared storage

This chapter shows sample RAID and LUN configurations with unshared storage, and then provides worksheets for planning your own shared storage installation. Topics are

- Paths to LUNs
- Sample unshared storage configurations
- Unshared storage planning worksheets

Dual SPs and paths to LUNs

With two SPs, you have two routes to LUNs. The SP on which you bind a disk is the default owner of the resulting LUN. The route through the SP that owns a unit determines the primary route to that LUN. The route through the other SP is the secondary route to the LUN.

With a single server, you can bind disks to each SP to share the I/O load equally. With dual servers, each server has one primary SP; the SP that binds a LUN determines the server that's the primary owner of the LUN.

Unshared storage system configurations

This section explains the different options available for connecting storage systems to servers. As needs change, you may want to change a configuration. You can do so without changing your LUN configuration or losing user data.

There are three types of configuration:

- Single-server configurations
- Dual-server configurations
- Configurations with hubs

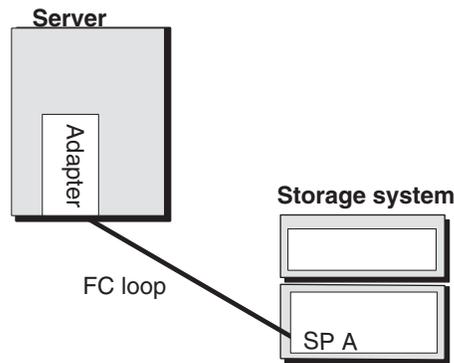
Single-server configurations

There are two single-server storage-system configurations: basic and dual-adapter/dual-SP. The following table and pictures show the features of each.

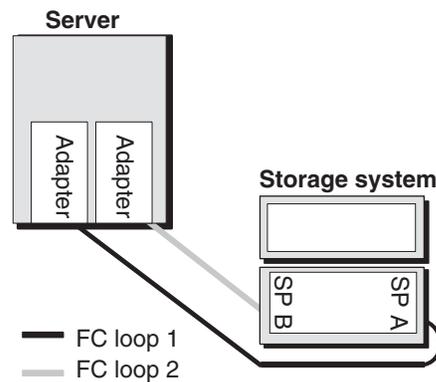
Storage-system configurations for a single server

Configuration type	Features
Basic One server One adapter One cable One SP	With a RAID group of any level other than 0, applications can continue after failure of any disk, but cannot continue after failure of an adapter, loop, or SP.
Dual-adapter/dual-SP One server Two adapters Two cables Two SPs	Provides highest availability and best storage-system performance of the single-server configurations; recommended for high-availability servers. With a RAID group of any level other than 0, applications can continue after any disk fails. If one adapter, loop, or SP fails, the software or system operator can transfer control to the other adapter.

Basic configuration



Dual-adapter/dual-SP configuration



Dual-server configurations

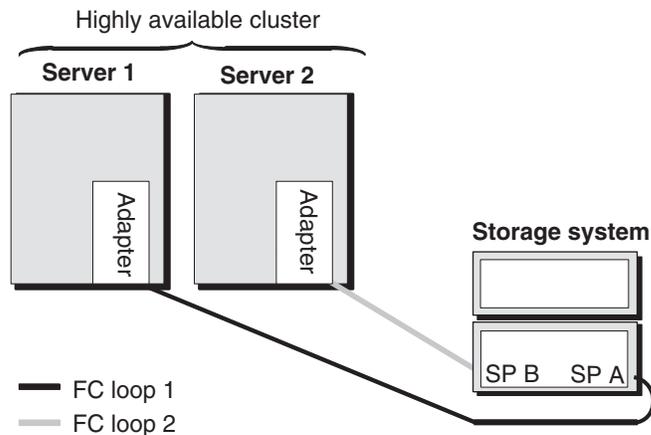
Dual-server configurations offer higher availability: if one server fails, another server can take over and run applications. Sharing storage systems between servers can lower costs. When two or more servers connect to the same storage system, the servers should run cluster software to retain possession of their disks.

There are two dual-server configurations: single-initiator/dual-loop and dual-initiator/dual-loop. The following table and pictures show the features of each.

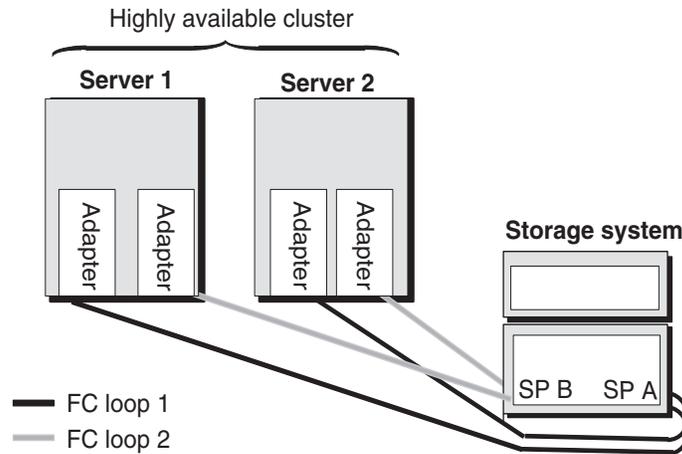
Dual-server configurations

Configuration type	Features
Single-initiator/dual-loop Two servers Two adapters (one per server) Two cables (one per server) Two SPs	Resembles two basic configurations using the same storage system, and provides some high availability for two servers. Each server and its applications can continue after failure of any disk. A server using a failed adapter or SP cannot continue after failure, but the other server can continue and run the first server's applications. Some operating systems support failover software that can automatically direct one server to take over the other's disks if the other fails. Or the system operator can transfer the disks manually.
Dual-initiator/dual-loop Two servers Four adapters (two per server) Four cables Two SPs	Two servers share two Fibre Channel loops. This provides the highest availability and best storage-system performance for dual-server configurations. With a RAID group of any level other than 0, applications can continue after any disk fails. If a server, adapter, cable, or SP fails, the system operator can transfer the affected disks to the other server, and then restart applications. Or the transfer and restart can occur automatically with failover software.

Single-initiator/dual-loop configuration



Dual-initiator/dual-loop configuration



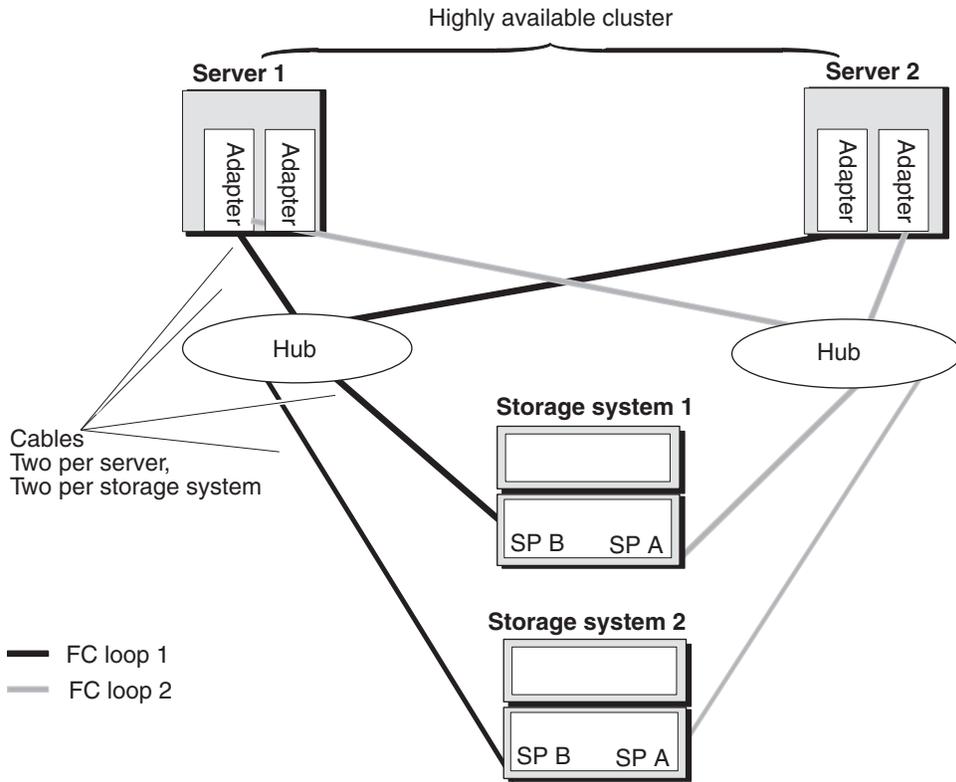
Configurations with hubs

A hub adds flexibility and availability to a Fibre Channel site. Essentially, a hub lets you expand any of the previous configurations to include multiple servers and storage systems. In any of these configurations, you can have one or more host-bus adapters in the servers, one or more hubs, and one or two SPs in each storage system. Each hub has nine ports as shown on page 1-6.

In any hub configuration, you cable the server(s) to the hub, and the hub to the storage system(s). To add a storage system, cable it to the hub; all servers connected to the hub then have access to that storage system. To add a server, cable it to the hub; the server has access to all storage systems connected to the hub. This eliminates the vulnerability of the daisy chain, since any storage system can be shut down without affecting the others.

The distance between nodes (servers, hubs, and storage systems) depends on the type of cable between them. An adapter that supports optical cable can exploit optical cable's speed and distance features; an adapter that does not support optical cables must use copper cables. A hub supports copper cables and — with MIAs — optical cables.

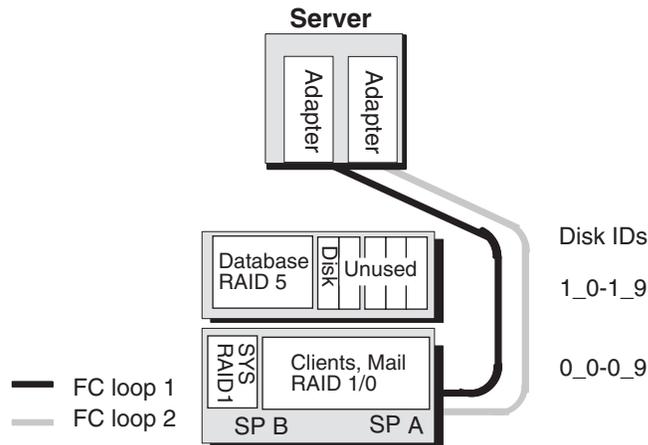
Sample configuration with hubs and high-availability options



Sample unshared storage configurations

This section shows disks in three unshared storage-system configurations — single-server dual-loop without a hub, dual-initiator/dual-loop without a hub, and multiple server dual-loop with two hubs.

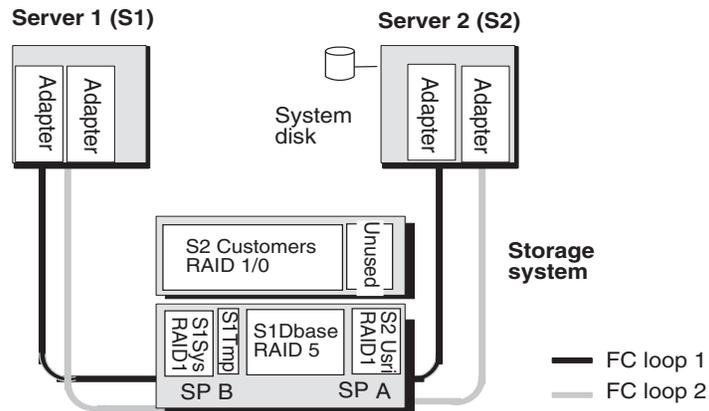
Single server example without hub



The storage system disk IDs and LUNs are as follows. The LUN capacities shown assume 18-Gbyte disks.

LUNs - SP A and SP B, 234 Gbytes		
<u>Disks RAID type. storage type</u>		
0_0, 0_1	– RAID 1, System disk,	18 Gbytes
0_2-0_9	– RAID 5 (8 disks), Clients and Mail,	126 Gbytes
1_0-1_4	– RAID 5, Database,	72 Gbytes
1_5	– Disk, Temporary storage,	18 Gbytes

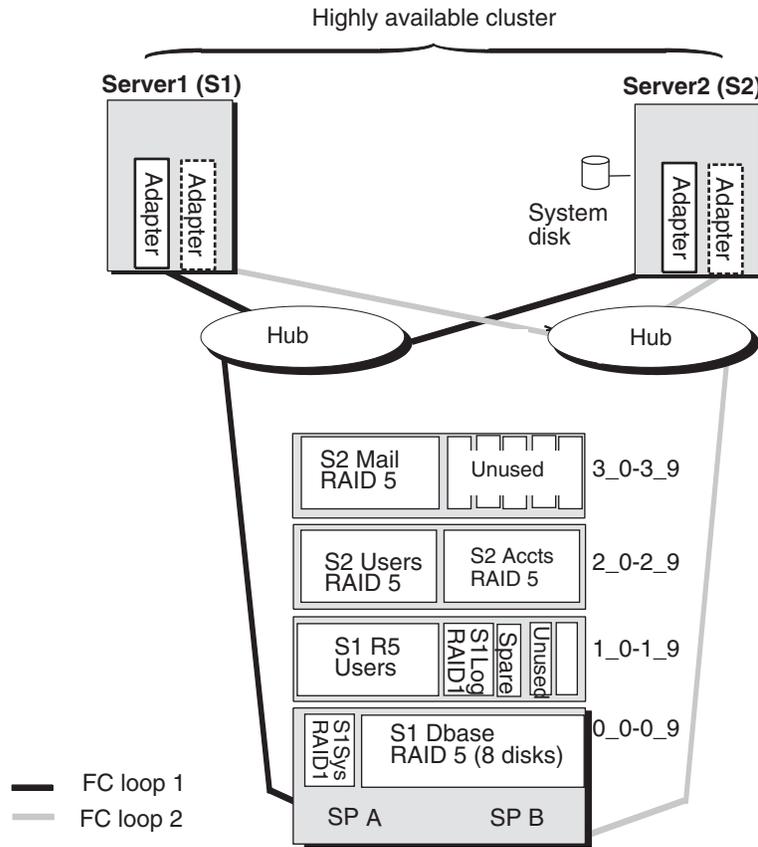
Dual-server example without hub



If each disk holds 18 Gbytes, then the storage-system chassis provides Server 1 with 108 Gbytes of disk storage, 90 Gbytes highly available; it provides Server 2 with 90 Gbytes of storage, all highly available. Each server has its own SP, which controls that server’s LUNs; those LUNs remain primary to that server. The LUNs are as follows. The LUN capacities shown assume 18-Gbyte disks.

Server1 LUNs (S1) - SP B, 126 Gbytes	Server2 LUNs (S2) - SP A, 126 Gbytes
<u>Disks RAID type. storage type. capacity</u>	<u>Disks RAID type. storage type. capacity</u>
0_0, 0_1 – RAID 1, System disk, 18 Gbytes	1_0-1_7 – RAID 5 (8 disks), Customer Accounts, 126 Gbytes
0_2 – Disk, Temporary storage, 18 Gbytes	
0_3-0_7 – RAID 5, Database, 72 Gbytes	
0_8-0_9 – RAID 1, Users, 18 Gbytes	

Multiple server example with hubs



The storage system disk IDs and servers' Storage Group LUNs are as follows. The LUN capacities shown assume 18-Gbyte disks.

Server1 LUNs (S1) - SP A, 234 Gbytes	Server2 LUNs (S2) - SP B, 216 Gbytes
<u>Disks RAID type, storage type</u>	<u>Disks RAID type, storage type</u>
0-0, 0_1 – RAID 1, System disk, 18 Gbytes	2_0-2_4 – RAID 5, Users, 72 Gbytes
0_2-0_9 – RAID 5 (8 disks), Database, 126 Gbytes	2_5-2_9 – RAID 5, Accounts, 72 Gbytes
1_0-1_4 – RAID 5, Users, 72 Gbytes	3_0-3_4 – RAID 5, Mail, 72 Gbytes
1_5,1_6 – RAID 1, Log files, 18 Gbytes	
1_7 – Hot spare (usable in any server's LUN)	

With a dual-loop configuration, for highest availability, all the servers should run the same operating system and participate in a cluster. The cluster environment will prevent one server from inadvertently overwriting the data owned by another server.

Planning applications and LUNs - Unshared storage

This section helps you plan your unshared storage use — applications you want to run and the LUNs that will hold them. The worksheets to help you do this include

- Application and file system planning worksheet - lets you outline your storage needs.
- LUN planning worksheet - lets you decide on the disks that will compose the LUNs.
- 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 shared storage system.

Sample file system, Storage Group, and LUN worksheets appear later in this chapter.

configuration. For example, for drive T, assign the volume ID RAID5_T. The volume label will then identify the drive letter.

The **RAID type of LUN** is the RAID group type you want for this partition, file system, or logical volume. 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.

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. Some systems allow only eight LUNs (numbers 0 through 7). For an operating system with this restriction, if you want a hot spare, assign the hot spare an ID above 7; for example, 8 or 9. The operating system never accesses a hot spare, so the ID is irrelevant to it.

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

For **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.

If this storage system will be used by two servers, provide a copy of this worksheet to the other server. This is particularly important where one server may take over the other's LUNs. If a LUN will be shared, on the Notes section of the LUN details worksheet, write "Primary to *server-name*" or "Secondary to *server-name*."

LUN planning worksheet

Use one of the following worksheets (Rackmount or Deskside) to select the disks that will make up the LUNs. Depending on model, a full-fibre rackmount storage system can include up to 120 disks, numbered 0 through 119, left to right from the bottom up. A 30-slot disk system with SCSI disks has a 30-disk maximum and different disk IDs.

Again depending on model, a deskside storage system can hold ten, 20, or 30 disks.

LUN planning worksheet - Rackmount

Full-fibre storage system

11_0	11_1	11_2	11_3	11_4	11_5	11_6	11_7	11_8	11_9
------	------	------	------	------	------	------	------	------	------

10_0	10_1	10_2	10_3	10_4	10_5	10_6	10_7	10_8	10_9
------	------	------	------	------	------	------	------	------	------

9_0	9_1	9_2	9_3	9_4	9_5	9_6	9_7	9_8	9_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

8_0	8_1	8_2	8_3	8_4	8_5	8_6	8_7	8_8	8_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

7_0	7_1	7_2	7_3	7_4	7_5	7_6	7_7	7_8	7_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

6_0	6_1	6_2	6_3	6_4	6_5	6_6	6_7	6_8	6_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

5_0	5_1	5_2	5_3	5_4	5_5	5_6	5_7	5_8	5_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

4_0	4_1	4_2	4_3	4_4	4_5	4_6	4_7	4_8	4_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

3_0	3_1	3_2	3_3	3_4	3_5	3_6	3_7	3_8	3_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

2_0	2_1	2_2	2_3	2_4	2_5	2_6	2_7	2_8	2_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

1_0	1_1	1_2	1_3	1_4	1_5	1_6	1_7	1_8	1_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

0_0	0_1	0_2	0_3	0_4	0_5	0_6	0_7	0_8	0_9
-----	-----	-----	-----	-----	-----	-----	-----	-----	-----

30-slot SCSI-disk storage system

A0	B0	C0	D0	E0	A3	B3	C3	D3	E3
----	----	----	----	----	----	----	----	----	----

A1	B1	C1	D1	E1	A4	B4	C4	D4	E4
----	----	----	----	----	----	----	----	----	----

A2	B2	C2	D2	E2	A5	B5	C5	D5	E5
----	----	----	----	----	----	----	----	----	----

Storage system number _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN planning worksheet - Deskside

Full-fibre storage system

0_0	1_0	2_0
0_1	1_1	2_1
0_2	1_2	2_2
0_3	1_3	2_3
0_4	1_4	2_4
0_5	1_5	2_5
0_6	1_6	2_6
0_7	1_7	2_7
0_8	1_8	2_8
0_9	1_9	2_9

30-slot SCSI disk storage system

A2	A1	A0
B2	B1	B0
C2	C1	C0
D2	D1	D0
E2	E1	E0
A5	A4	A3
B5	B4	B3
C5	C4	C3
D5	D4	D3
E5	E4	E3

Storage system number _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

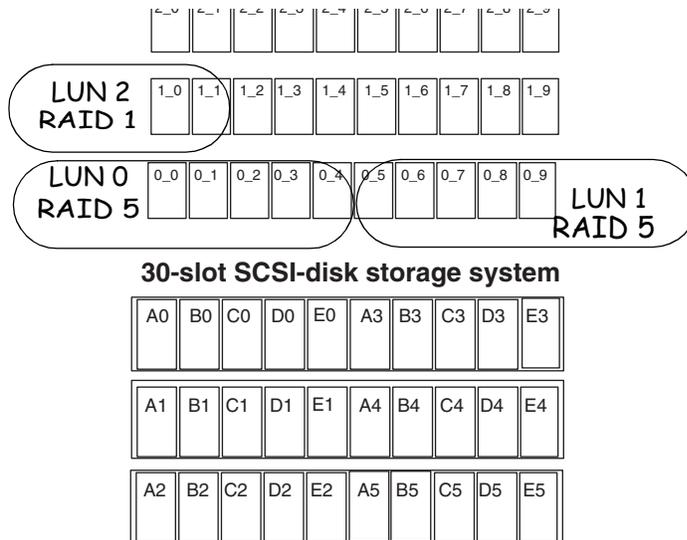
LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

LUN number _____ RAID type ____ Cap. (Gb) _____ Disk IDs _____

A sample LUN worksheet follows.



Storage system number	<u>1</u>						
LUN number	<u>0</u>	RAID type	<u>5</u>	Cap. (Gb)	<u>72</u>	Disk IDs	<u>0_0, 0_1, 0_2, 0_3, 0_4</u>
LUN number	<u>1</u>	RAID type	<u>5</u>	Cap. (Gb)	<u>72</u>	Disk IDs	<u>0_5, 0_6, 0_7, 0_8, 0_9</u>
LUN number	<u>2</u>	RAID type	___	Cap. (Gb)	<u>18</u>	Disk IDs	<u>1_0, 1_1</u>
LUN number	___	RAID type	___	Cap. (Gb)	___	Disk IDs	___

Completing the LUN 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. This form is *enclosure_diskID*, where *enclosure* is the enclosure number (the bottom one is 0, above it 1, and so on) and *diskID* is the disk position (left is 0, next is 1, and so on).

IMPORTANT On a full-fibre system, none of the disks 0_0 through 0_8 may be used as a hot spare. On a SCSI-disk system, none of the disks A0, B0, C0, or A3 may be used as a hot spare.

Next, complete as many of the LUN sections as needed for each storage system. Copy the (blank) worksheet as needed for all LUNs in each storage system. A storage system is any group of enclosures connected to a DPE; a full-fibre system can include up to 11 DAE enclosures for a total of 120 disks. If a Storage Group

LUN Details worksheet

Use the following LUN details worksheet to plan the individual LUNs. Complete as many of these as needed for all LUNs.

LUN Details Worksheet

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

Storage-system number or name: _____

Storage-system configuration

Shared storage

Unshared storage:

Basic

Dual-adapter/dual-SP

Single-initiator/dual-loop

Dual-initiator/dual-loop

Hub, single loop

Hub, dual loop

SP FC-AL address ID (unshared only):

SP A: _____ SP B: _____

SP memory (Mbytes):

SP A: _____ SP B: _____

Use for caching

Read cache size: _____ MB Write cache size: _____ MB Cache page size: _____ KB

Use for RAID 3

LUN ID: _____

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

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: _____

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

LUN ID: _____

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

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: _____

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

LUN ID: _____

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

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: _____

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

LUN ID: _____

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

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: _____

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

Completing the LUN details worksheet

Complete the header portion of the worksheet for each storage system as described below. Copy the blank worksheet as needed. Sample completed LUN worksheets appear later.

Storage-system entries

Storage-system configuration. Specify Unshared. Then specify the unshared configuration; for example , single-initiator, dual-loop.

For any multiple-server configuration, each server will need cluster software.

SP FC-AL address ID. For unshared storage, which uses FC-AL addressing, each SP (and each other node) on a Fibre Channel loop must have a unique FC-AL address ID. You set the SP FC-AL address ID using switches on the back panel of the SP. The valid FC-AL address ID range is a number 0 through 125 decimal, which is 0 through 7D hexadecimal. For any number above 9, we suggest hexadecimal, since the switches are marked in hexadecimal.

If you have two FC-AL loops, we suggest a unique FC-AL address ID for each SP on both loops.

SP memory. Enter the amount of memory each SP has. If a storage system has two SPs, they will generally have the same amount of memory. You can allocate this memory to storage-system caching or RAID 3 use.

Use memory for caching. You can use SP memory for read/write caching or RAID 3. (Using both caching and RAID 3 in the same storage system not recommended.) If you choose caching, check the box and continue to the next step; for RAID 3, skip to the RAID Group ID entry.

Read cache size. If you want a read cache, it should generally be about one third of the total available cache memory.

Write cache size. The write cache should be two thirds of the total available. Some memory is required for system overhead, so you cannot determine a precise figure at this time. For example, for 256 Mbytes of total memory, you might have 240 Mbytes available, and you would specify 80 Mbytes for the read cache and 160 Mbytes for the write cache.

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

For a general-purpose file server — 8 Kbytes

For a database application — 2 or 4 Kbytes

The ideal cache page size depends on the operating system and application.

Use memory for 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. Some systems allow only eight LUNs (numbers 0 through 7). For an operating system with this restriction, if you want a hot spare, assign the hot spare an ID above 7; for example, 8 or 9. The operating system never accesses a hot spare, so the ID is irrelevant to it.

RAID Group ID. This 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).

For **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 18-Gbyte disks holds 72 Gbytes;
- An eight-disk RAID 1/0 group of 18-Gbyte disks also holds 72Gbytes;
- A RAID 1 mirrored pair of 18-Gbyte disks holds 18 Gbytes; and
- An individual disk of a 9-Gbyte disk also holds 18 Gbytes.

Each disk in the RAID group must have the same capacity; otherwise, you will waste disk storage space.

LUN Size. Enter the user-available capacity in gigabytes (Gbytes) of the LUN. You can make this the same size as the RAID group, above. 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 ID(s) 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 the DPE (enclosure 0 disks 2 through 6, enter 0_2, 0_3, 0_4, 0_5, and 0_6.

SP. 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.

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 4-16), you can specify whether you want caching — 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.

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. Enter the name of each server that will be able to use the LUN. Normally, you need to restrict access by establishing SP ownership of LUNs when you bind them.

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 this storage system's chassis will be shared with another server, and the other server is the primary owner of this disk, write "secondary." (If the storage system will be used by two servers, we suggest you complete one of these worksheets for each server.)

What next?

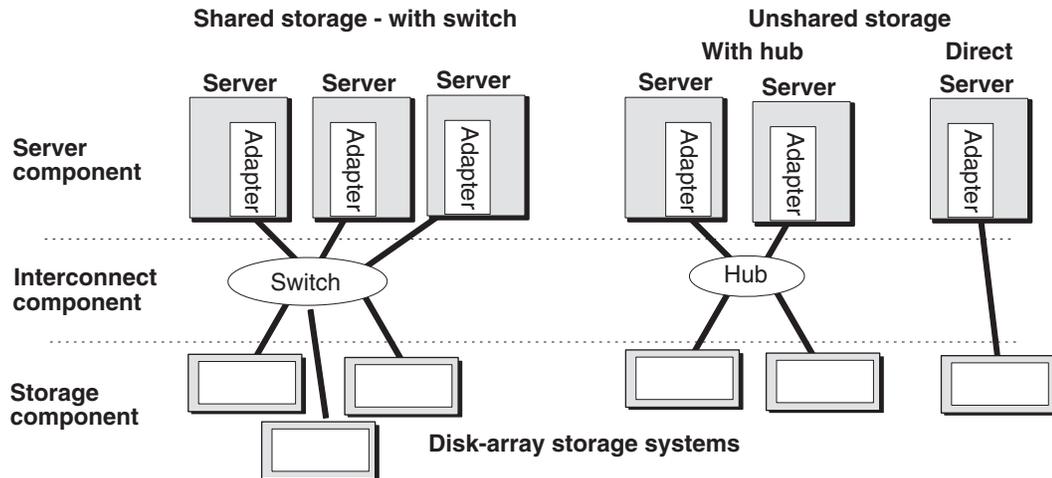
This chapter outlined the planning tasks for unshared storage systems. If you have completed the worksheets to your satisfaction, you are ready to learn about the hardware needed for these systems as explained in Chapter 5.



5

Storage-system hardware

The products described in this chapter compose the storage component of storage systems. The storage systems attach to the server and the interconnect components described in Chapter 1.



Topics in this chapter are

- Hardware for shared storage
- Hardware for unshared storage
- Planning your hardware components
- Component data sheets
- Cabinets for rackmount enclosures
- Cable and configuration guidelines
- Component planning worksheets

Hardware for shared storage

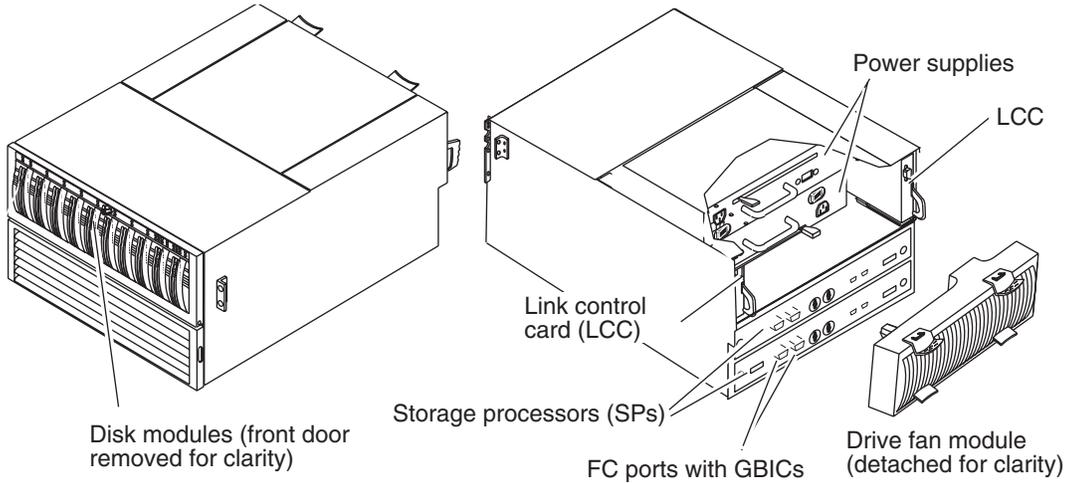
The primary hardware component for shared storage is a ten-slot Disk-array Processor Enclosure (DPE) with two storage processors (SP). The DPE can support up to 11 separate 10-slot enclosures called Disk Array Enclosures (DAEs) for a total of 120 disks. Shared storage requires two SPs and LIC of a specific model.

A DPE with a DAE is available as a deskside system, but with a capacity of 20 disks this cannot provide the expandability and total storage capacity needed for a SAN (storage area network). So this section does not cover the deskside version.

Storage hardware — rackmount DPE-based storage systems

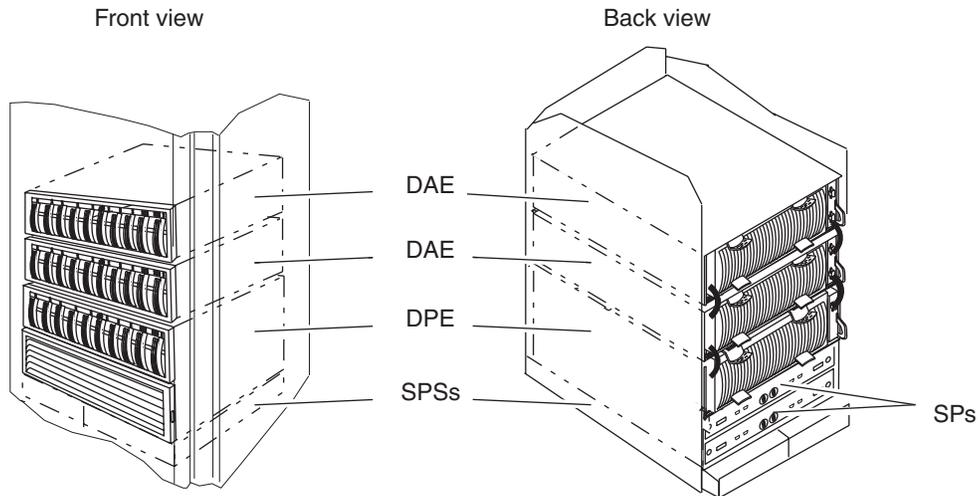
The DPE rackmount enclosure is a sheet-metal housing with a front door, a midplane, and slots for the storage processors (SPs), link control cards (LCCs), disk modules, power supplies, and fan packs. All components can be replaced under power. The DPE rackmount model looks like the following figure.

DPE storage-system components – rackmount model



A separate standby power supply (SPS) is required to support write caching. All the shared storage components — rackmount DPE, DAEs, SPSs, and cabinet — are shown in the following figure.

Rackmount system with DPE and DAEs

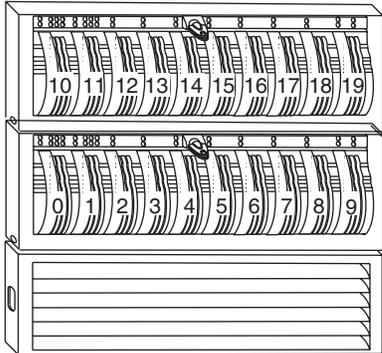


Disks

The disks — available in differing capacities — fit into slots in the enclosure. Each module has a unique ID that you use when binding or

monitoring its operation. The ID is derived from the enclosure address (always 0 for the DPE, settable on a DAE) and the disk module slot numbers.

Disk modules and module IDs — rackmount DPE-based system

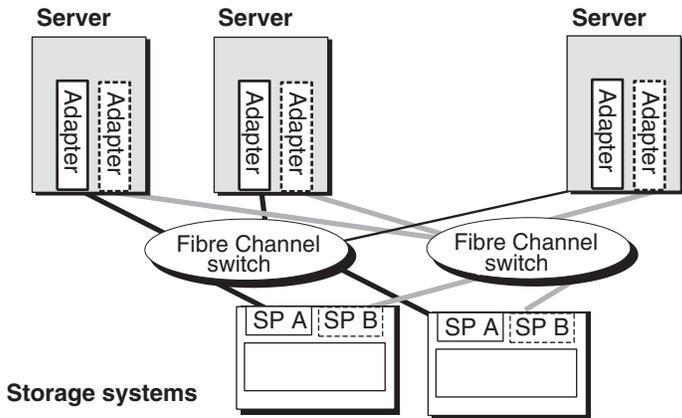


Storage processor (SP)

The SP provides the intelligence of the storage system. Using its own operating system (called Licensed Internal Code), the SP processes the data written to or read from the disk modules, and monitors the modules themselves. An SP consists of a printed-circuit board with memory modules (DIMMs), and status lights.

For high availability, a storage system can support a second SP. The second SP provides a second route to a storage system and also lets the storage system use write caching for enhanced write performance. Two SPs are required for shared storage.

Shared storage systems



There are more examples of shared storage in Chapter 3.

Hardware for unshared storage

Unshared storage systems are less costly and less complex than shared storage systems. They offer many shared storage system features; for example, you can use multiple unshared storage systems with multiple servers. However, with multiple servers, unshared storage offers less flexibility and security than shared storage, since any user with write access to privileged server files can enable access to any storage system.

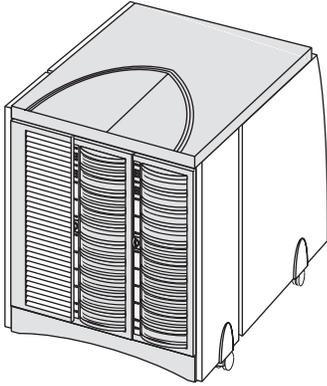
Types of storage system for unshared storage

For unshared storage, there are four types of storage system, each using the FC-AL protocol. Each type is available in a rackmount or desktside (office) version.

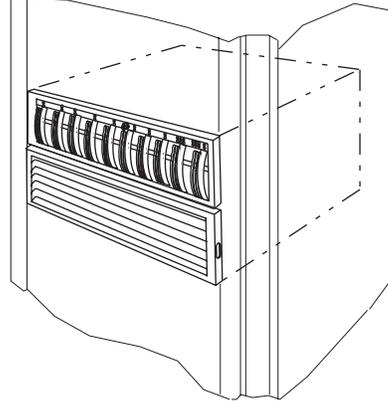
- **Disk-array Processor Enclosure (DPE)** storage systems. A DPE is a 10-slot enclosure with hardware RAID features provided by one or two storage processors (SPs). In addition to its own disks, a DPE can support up to 110 additional disks in 10-slot Disk Array Enclosures (DAEs) for a total of 120 disks. This is the same kind of storage system used for shared storage, but it uses a different storage processor (SP).
- **Intelligent Disk Array Enclosure (iDAE)**. An iDAE, like a DPE, has SPs and thus all the features of a DPE, but is thinner and has a limit of 30 disks.
- **Disk Array Enclosure (DAE)**. A DAE does not have SPs. A DAE can connect to a DPE or an iDAE, or you can use it without SPs. A DAE used without an SP does not inherently include RAID, but can operate as a RAID device using software running on the server system. Such a DAE is also known as **Just a Box of Disks**, or JBOD.
- **30-slot SCSI-disk storage systems**. Like the DPE, these offer RAID features provided by one or two SPs. However, they use SCSI, not Fibre Channel, disks. Each has space for 30 disks.

Disk-array processor enclosure (DPE)

Deskside DPE with DAE

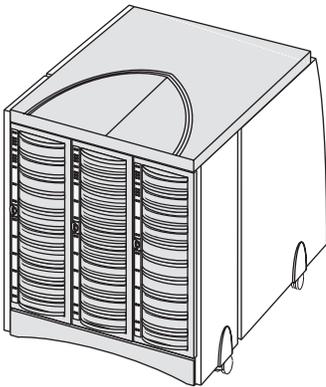


Rackmount DPE, one enclosure, supports up to 11 DAEs

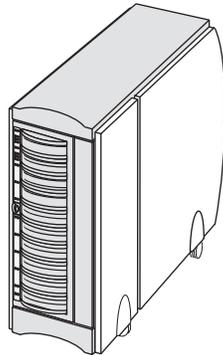


Intelligent disk-array enclosure (IDAE)

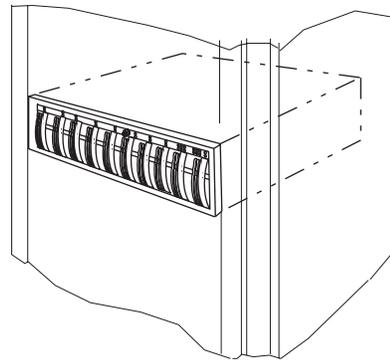
30-slot deskside



10-slot deskside

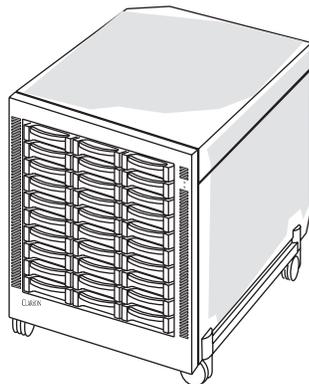


Rackmount

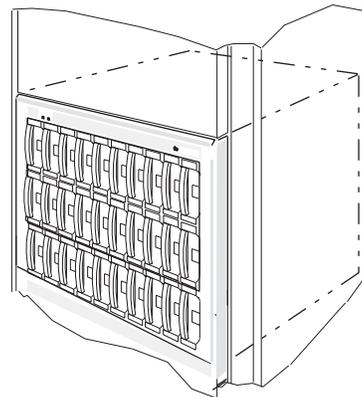


30-slot SCSI-disk storage system

Deskside

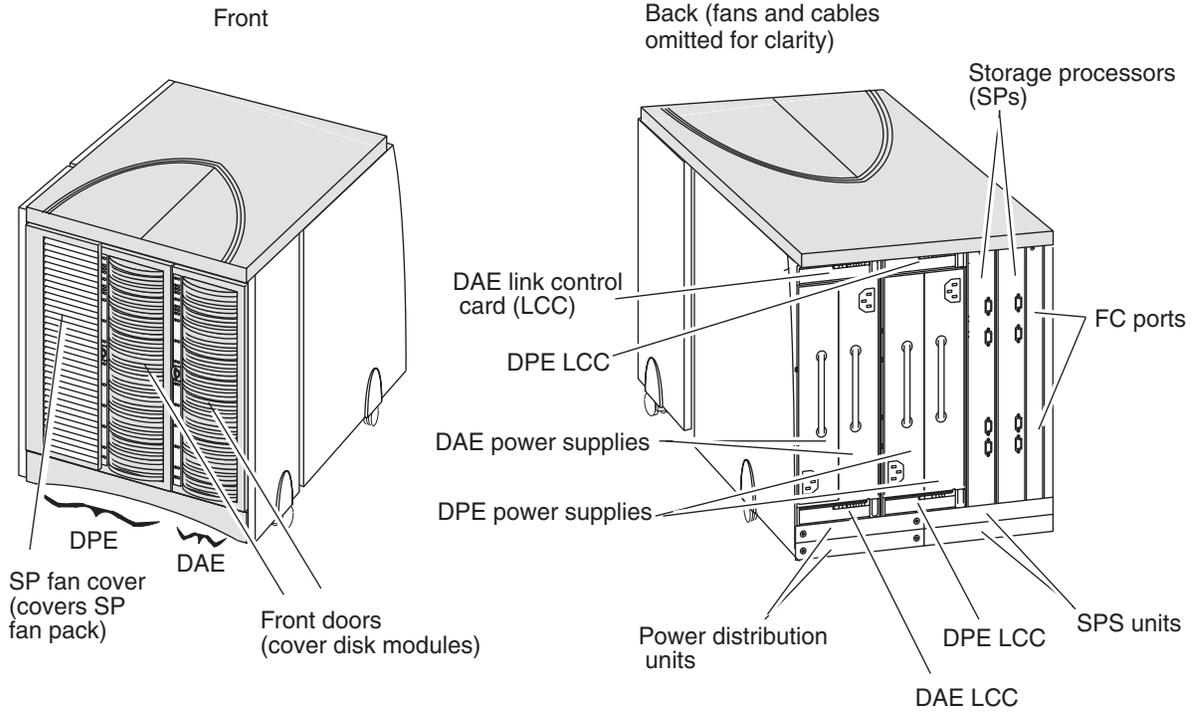


Rackmount



The following figure shows some components of a desktide DPE. Components for the other types are similar.

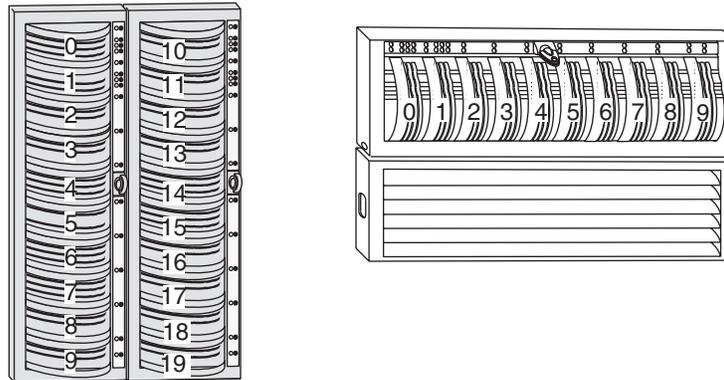
DPE components – desktide model



Disks

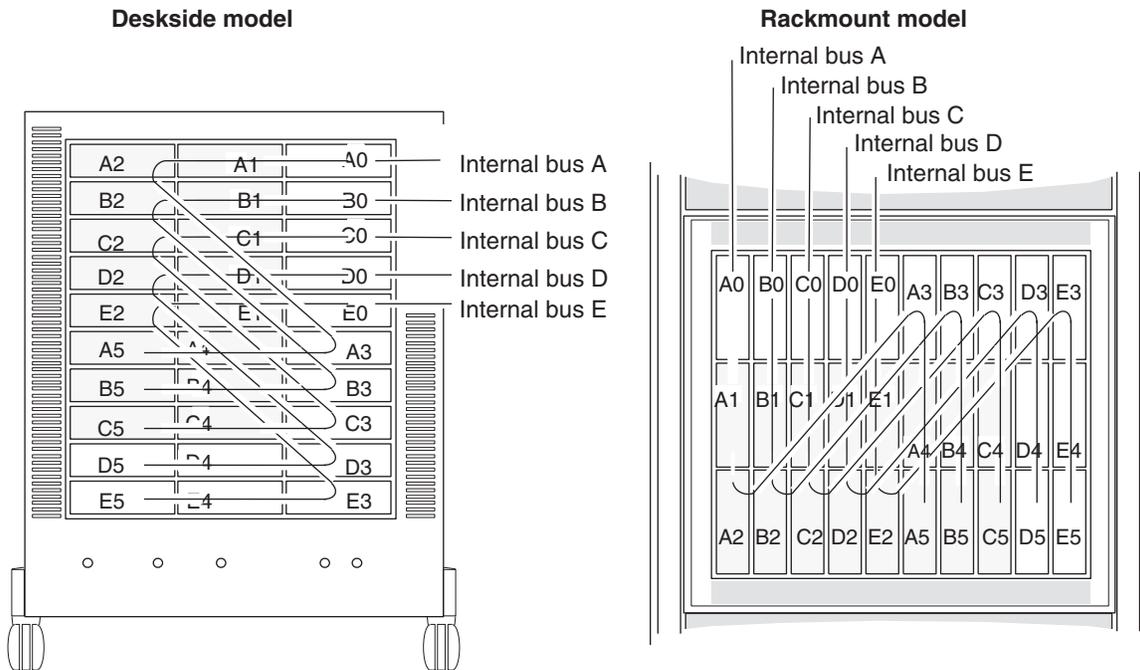
The disks — available in differing capacities — fit into slots in the enclosure. Each disk has a unique ID that you use when binding it or monitoring its operation. The ID is the enclosure address (always 0 for the DPE, settable on a DAE) and the disk slot number.

Disks and disk IDs in a full-fibre system



For a 30-slot system with SCSI disks, the disk IDs are based on the SCSI internal bus and the number of the disk on that bus, as follows:

Disks and disk IDs in a 30-slot SCSI disk system

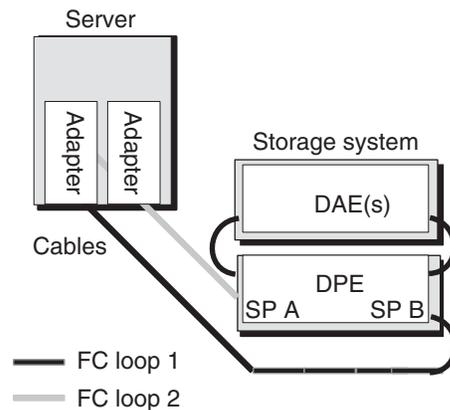


Storage processor (SP)

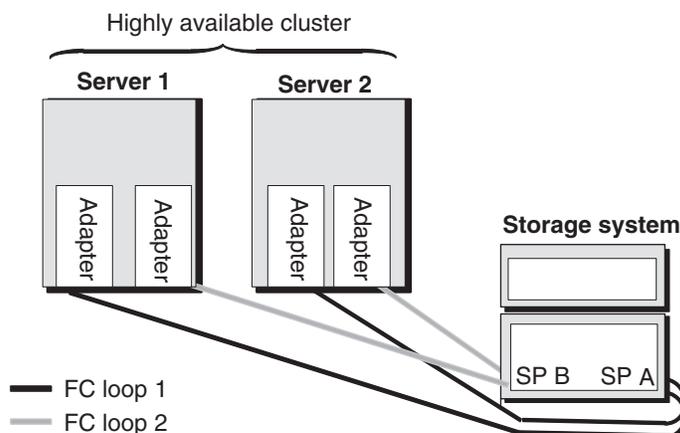
The SP provides the intelligence of the storage system. Using its own operating system (called Licensed Internal Code), the SP processes the data written to or read from the disk modules, and monitors the modules themselves. An SP consists of a printed-circuit board with memory modules (DIMMs), status lights, and switches for setting FC-AL addresses.

For high availability, a storage system can support a second SP. A second SP provides a second route to a storage system, so both SPs can connect to the same server or two different servers, as follows.

Storage system with two SPs connected to the same server



Storage system with two SPs connected to different servers



Either SP can control any LUN in the storage system, but only one SP at a time can control a LUN. If one SP cannot access a LUN it controlled (because of a failure), you can transfer control of the LUN to the other SP, manually or via software.

Storage-system caching provides significant performance enhancement. Read caching is available with one or two SPs. Mirrored write caching, particularly helpful with RAID 5 I/O, requires two SPs (to mirror one another, for cache integrity) and a Standby Power Supply (SPS) to enable the SPs to write their cached data to disk if power fails.

Planning your hardware components

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

For shared storage, you must use a DPE rackmount system with two SPs and high-availability options. We assume you have some idea of how many servers, adapters, hubs, storage systems, and SPs you want. Skip to the component data sheets following.

For unshared storage, you can use one or two SPs and you can choose among storage system configurations. This section assumes you have examined the configurations shown starting on page 4-1 and have some idea of how many servers, adapters, hubs, storage systems, and SPs you want. It ends with blank worksheets and sample worksheets.

Configuration tradeoffs - shared storage

The hardware configuration required for shared storage is very specific: two host-bus adapters in each attached server, two Fibre Channel switches, and two SPs per storage system. Choices you can make with shared storage systems include the number of storage systems (up to 15 are allowed), and for each storage system the cache configuration (maximum or minimum), and one or two standby power supplies (SPS units).

The number of storage systems in the SAN depends on the servers' processing demands. For each system, the larger cache improves write performance for very large processing loads; the redundant SPS lets write caching continue if one SPS fails.

Configuration tradeoffs - unshared storage

For each storage-system enclosure, you have two important areas of choice: rackmount or deskside model, and high-availability options.

Generally, rackmount systems are more versatile; you can add capacity in a cabinet without consuming more floor space. However, rackmount systems require additional hardware, such as cabinets and mounting rails, and someone must connect power cords and cables within them. For large storage requirements, rackmount systems may be more economical than deskside systems. Deskside systems are more convenient; they ship with all internal cabling in place and require only ac power and connection to the servers.

For high availability, there are many variations. The most important high-availability features are a second SP/LCC pair, second power supply, and standby power supply (SPS). The second SP/LCC and SPS let you use write caching to enhance performance; the second SP provides continuous access to storage-system disks if one SP or LCC fails. Another high-availability option is a redundant SPS.

Yet another option, for a deskside system, is a second power distribution unit (PDU), which lets you route ac power from an independent source. Used this way, the second PDU protects against failure in one of the two ac power sources. With a rackmount system, you can acquire a cabinet with one or two ac inlet cords. The second inlet cord, connected to a second ac power source, provides the same advantage for all storage systems in the cabinet as the second PDU in the deskside storage system.

For deskside systems, the optional high-availability hardware fits into the deskside cabinet. Deskside high-availability options are as follows.

High-availability options, desktide unshared storage

Desktide system type	HA level	PDUs	SPs	LCCs	Power supplies	Disks	SPS units
DPE	Minimum	1	1	1 DPE 1 DAE	1 DPE 1 DAE	5 (without write cache)	0 (without write cache) 1 (write cache)
	Maximum	2	2	4 (2 DPE 2 DAE)	4 (2 DPE 2 DAE)	10 (write cache or RAID 3)	2
iDAE	Minimum	1	1	n/a (10-slot) 2 (30-slot)	1	3 (without write cache) 5 (write cache or RAID 3)	0 (without write cache) 1 (write cache)
	Maximum	2	2	n/a (10-slot) 4 (30-slot)	2 (10-slot) 6 (30-slot)	5 (write cache or RAID 3)	2
DAE only	Minimum	1	n/a	1	1	No minimum	n/a
	Maximum	2	n/a	2	2	No minimum	n/a
30-slot SCSI-disk	Minimum	n/a	1	n/a	1 (called a VSC)	5 (write cache or RAID 3)	0 (without write cache) 1 (write cache)
	Maximum	n/a	2	n/a	3 (called VSCs)	5 (write cache or RAID 3)	1

For rackmount systems, the standby power supply or supplies (SPS or BBU) must be placed in a tray directly beneath the storage system. Typically, any hubs in the cabinet mount at the top or bottom of the cabinet. Rackmount options are as follows.

High-availability options, rackmount unshared storage

Rackmount system type	HA level	SPs	LCCs	Power supplies	Disks	SPS units
DPE	Minimum	1	1	1	5 (without write cache)	0 (without write cache) 1 (write cache)
	Maximum	2	2 (DPE) 22 (with 11 DAEs)	2 (DPE) 22 (with 11 DAEs)	10 (write cache or RAID 3)	2
iDAE	Minimum	1	n/a (10-slot) 4 (with two DAEs)	1	3 (without write cache) 5 (write cache or RAID 3)	0 (without write cache) 1 (write cache)
	Maximum	2	n/a	2	5 (write cache or RAID 3)	2
DAE only	Minimum	n/a	1	1	No minimum	n/a
	Maximum	n/a	2	2	No minimum	n/a
30-slot SCSI-disk	Minimum	1	n/a	1 (called a VSC)	5 (write cache or RAID 3)	0 (without write cache) 1 (write cache)
	Maximum	2	n/a	3 (called VSCs)	5 (write cache or RAID 3)	1

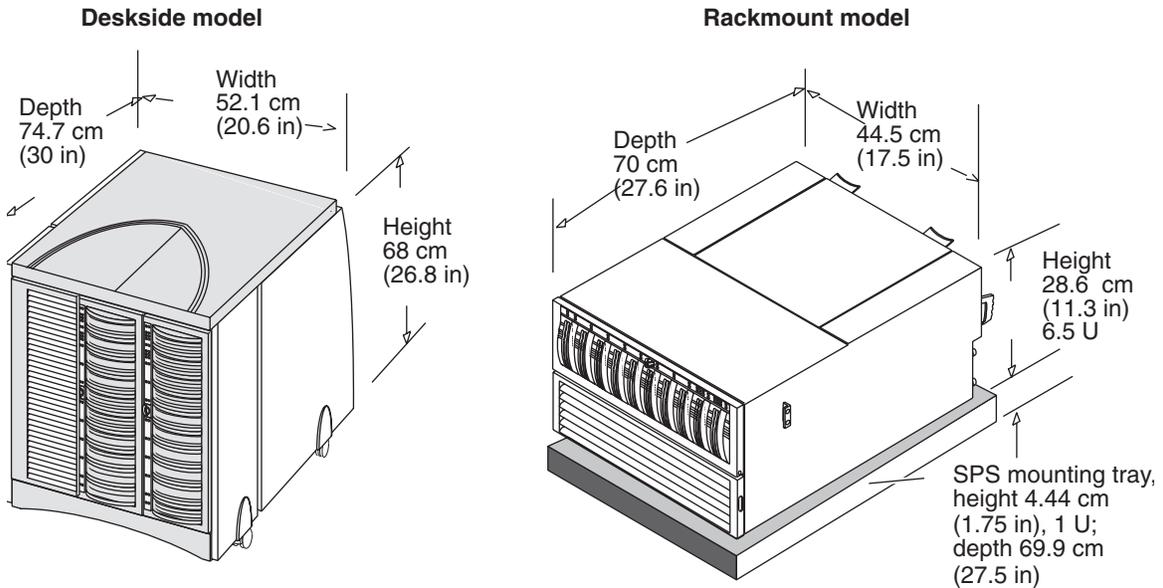
Hardware data sheets

The hardware data sheets shown in this section provide the plant requirements, including dimensions (footprint), weight, power requirements, and cooling needs, for DPE, iDAE, DAE, and 30-slot SCSI disk systems. Sections on cabinets and cables follow the data sheets.

DPE data sheet

For shared storage, a rackmount DPE and one or more rackmount DAEs are required. For unshared storage, you can use a rackmount or deskside DPE and DAE(s). The DPE dimensions and requirements are shown in the following figure.

DPE dimensions and requirements



Weight (without packaging)

Maximum (max disks, SPs, LCCs, PSs):
with 2 SPSs

Deskside

144 kg (316 lb)
165 kg (364 lb)

Rackmount

52 kg (115 lb)
74 kg (163 lb)

Power requirements

Voltage rating: 100 V ac to 240 V ac $-10\%/+15\%$, single-phase, 47 Hz to 63 Hz;
power supplies are auto-ranging
Current draw: At 100 v ac input – Deskside DPE/DAE: 12.0 A;
Rackmount DPE: 8.0 A max
SPS: 1.0 A max per unit during charge
Power consumption: Deskside DPE/DAE: 1200 VA; Rackmount DPE: 800 VA max
SPS: 100 VA per unit during charge

Power cables (single or dual)

ac inlet connector: IEC 320-C14 power inlet
Deskside power cord: USA: 1.8 m (6.0 ft): NEMA 6-15P plug
Outside USA: Specific to country

Operating environment

Temperature: 10°C to 40°C (50°F to 104°F)
Relative humidity: Noncondensing, 20% to 80%
Altitude: 40°C to 2,438 m (8,000 ft); 37°C to 3,050 m (10,000 ft)
Heat dissipation (max): Deskside DPE/DAE: 3931×10^3 J/hr (2730 BTU/hr) max estimated;
Rackmount DPE: 2520×10^3 J/hr (2390 BTU/hr) max estimated
Air flow: Front to back

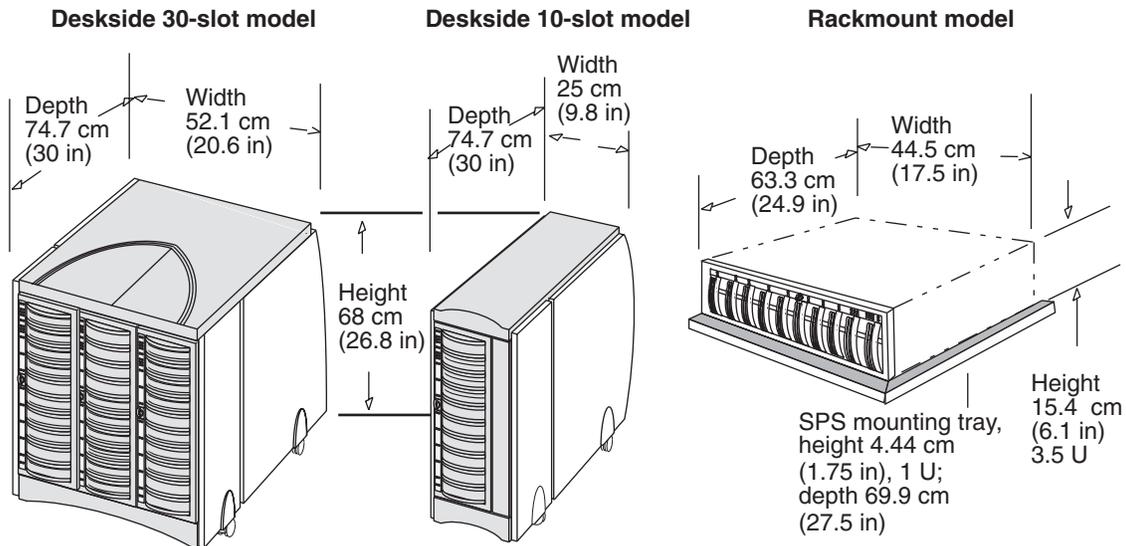
Service clearances

Front: 30.3 cm (1 ft)
Back: 60.6 cm (2 ft)

iDAE data sheet

You can use a rackmount or deskside DPE and DAE(s) for unshared storage. The iDAE dimensions and requirements are shown in the following figure.

Dimensions and requirements, iDAE



Weight (without packaging)

Maximum (max disks, SPS):

Deskside 30

144 kg (316 lb)

Deskside 10

60 kg (132 lb)

Rackmount

35.4 kg (78 lb)

Power requirements

Voltage rating:

100 V ac to 240 V ac +/- 10%, single-phase, 47 Hz to 63 Hz;
power supplies are auto-ranging

Current draw

At 100 v ac input – 30-slot 12.0 A; 10-slot: 4.0 A;
SPS: 1.0 A max per unit during charge

Power consumption:

30-slot: 1200; 10-slot: 400 VA;
SPS: 100 VA per unit during charge

Power cables (single or dual)

ac inlet connector:

IEC 320-C14 power inlet

Deskside power cord:

USA: 1.8 m (6.0 ft): NEMA 6-15P plug
Outside USA: Specific to country

Operating environment

Temperature:

10°C to 40°C (50° F to 104° F)

Relative humidity:

Noncondensing, 20% to 80%

Altitude:

40°C to 2,438 m (8,000 ft); 37°C to 3,050 m (10,000 ft)

Heat dissipation (max):

30-slot: 4,233 KJ/hr (4,020 BTU/hr)

10-slot: 1,411 KJ/hr (1,340 BTU/hr)

Air flow:

Front to back

Service clearances

Front:

30.3 cm (1 ft)

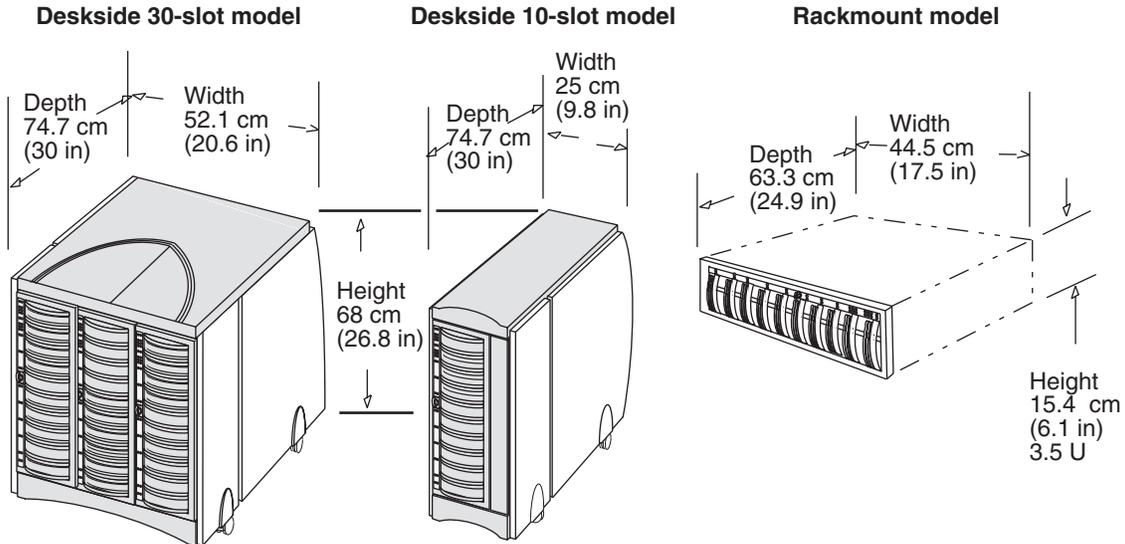
Back:

60.6 cm (2 ft)

DAE data sheet

The DAE storage-system dimensions and requirements are shown in the following figure.

Dimensions and requirements, DAE



Weight(without packaging)	Deskside 30	Deskside 10	Rackmount
Maximum configuration	143.6 kg (316 lb)	60 kg (132 lb)	35.4 kg (78 lb)

Power requirements

Voltage rating:	100 V ac to 240 V ac -10%/+15%, single-phase, 47 Hz to 63 Hz; power supplies are auto-ranging
Current draw:	At 100 V: 30-slot: 12 A; 10-slot 4 A max
Power consumption:	30-slot: 1200 VA; 10-slot 400 VA per supply max

Power cables (single or dual)

ac inlet connector:	IEC 320-C14 power inlet
Deskside power cord:	USA: 1.8 m (6.0 ft): NEMA 6-15P plug Outside USA: Specific to country

Operating environment

Temperature:	10°C to 40°C (50° F to 104° F)
Relative humidity:	Noncondensing, 20% to 80%
Altitude:	40°C to 2,438 m (8,000 ft); 37°C to 3,050 m (10,000 ft)
Heat dissipation (max):	30-slot: 4,233 KJ/hr (4,020 BTU/hr) 10-slot: 1,411 KJ/hr (1,340 BTU/hr)

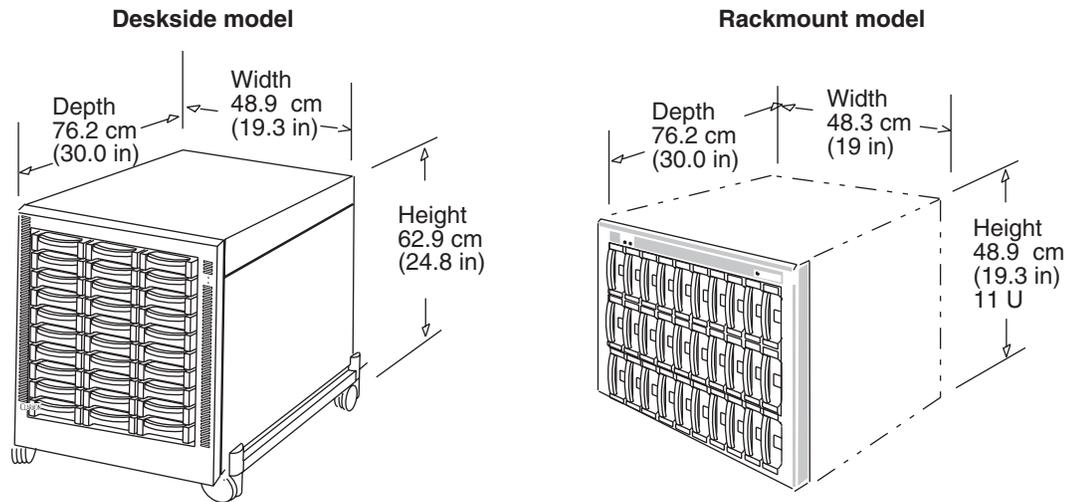
Service clearances

Front:	30.3 cm (1 ft)
Back:	60.6 cm (2 ft)

30-slot system with SCSI disks data sheet

The 30-slot disk-array storage system dimensions and requirements are shown in the following figure.

Dimensions and requirements, 30-slot SCSI-disk storage system



Weight(without packaging)

Maximum configuration:

Deskside

119 kg (261.5 lb)

Rackmount

98 kg (216 lb)

Power requirements

Voltage rating: 200 V ac to 240 V ac -10%/+15%, single-phase, 47 Hz to 63 Hz;
power supplies are auto-ranging

Current draw: At 5.0 maximum at 200 Vac input

Power consumption: 1000 VA maximum (apparent)

Power cables (single or dual)

ac inlet connector: IEC 320-C14 power inlet

Deskside power cord: USA: 1.8 m (6.0 ft): NEMA 6-15P plug

Outside USA:: Specific to country

Operating environment

Temperature: 10°C to 38°C (50° F to 100° F)

Relative humidity: Noncondensing, 20% to 80%

Altitude: 2,438 m (8,000 ft)

Heat dissipation (max): 300 BTU/hour

Air flow: Front to back

Service clearances

Front: 30.3 cm (1 ft)

Back: 60.6 cm (2 ft)

Cabinets for rackmount enclosures

Prewired 19-inch-wide cabinets, ready for installation, are available in the following dimensions to accept rackmount storage systems.

Vertical space	Exterior dimensions	Comments
173 cm or 68.25 in (39 NEMA units or U; one U is 1.75 in)	Height: 192 cm (75.3 in) Width: 65 cm (25.5 in) Depth: 87 cm (34.25 in) plus service clearances, which are 90 cm (3 ft), 30 cm front and 60 cm back Weight (empty): 134 kg (296 lb)	Accepts combinations of: DPEs at 6.5 U, iDAEs at 3.5 U, SPS units at 1 U, DAEs at 3.5 U each, 30-slot SCSI disk systems at 11 U Switches or hubs at 1 U Requires 200–240 volts ac. Plug options include L6–30 or L7–30 (domestic) and IEC 309 30 A (international). Each power strip has 12 IEC-320 CIS outlets. Filler panels of various sizes are available.

As an example, a rackmount storage system that supports 100 disk modules has the following requirements.

Category	Requirement
Vertical cabinet space in NEMA units (U, one U is 1.75 in)	Bottom to top: One SPS (1 U), one DPE (6.5 U), and nine DAEs (9*3.5 U equals 31.5 U) for a total of 39 U.
Weight	516 kg (1,137 lb) including the cabinet (134 kg), DPE (52 kg), SPS (11 kg), and nine DAEs (9 * 35.4 kg equals 319 kg).
Power	4,500 VA max, including the DPE (800 VA), SPS (100 VA), and nine DAEs (9 * 400 VA equals 3600 VA).
Cooling	15,484 KJ/hour (14,700 BTU/hr), including the DPE (2,520 KJ/hr), SPS (265 KJ/hour, estimated), and nine DAEs (9*1,411 KJ/hr equals 12,699 KJ/hr).

Cable and configuration guidelines

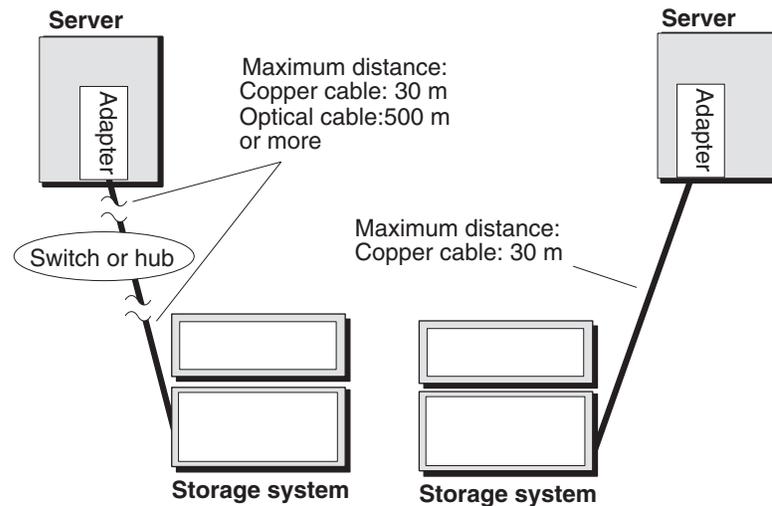
We recommend that all copper-interconnected nodes be connected to a common ground grid. The common grid is not needed for optical interconnections.

Copper cable allows up to 30 meters (99 feet) between nodes or hubs. Optical cable allows significantly longer distances. This is a major advantage of optical cable. However, you can use optical cable from a server only if the server's adapter supports optical cable; otherwise you must use copper. Not all adapters support optical cable.

To connect a DPE to a DAE, you must use copper cable, whose maximum length is 10 meters (33 ft). So, the distance between a DPE and the DAEs it controls cannot exceed 10 meters (33 ft).

The HBAs and SPs used with shared storage systems require optical cable, as does the switch between the HBAs and SPs. An optical GBIC interface converter, which is required for optical cable, is included with each SP and HBAs, but are not included with a switch; you must specifically order an optical GBIC for every switch port you will use.

The SPs used with unshared storage systems support copper cables and — with MIAs — optical cables. The hub itself supports copper, or with a MIA, optical. So you can use a copper cable or — with two MIAs per cable — optical cable between any hub and SP. For optical cable to work between an HBA and hub or SP, the HBA must support optical cable.



You can use any existing FDDI, multimode, 62.5 micron cable with good connections to attach servers, hubs, and storage systems. These cables must be dedicated to storage-system I/O.

Cable sizes — Optical

Length	Typical use
5 m (16.5 ft) or 10 m (33 ft)	Within one room, connecting servers to storage systems (adapter must support optical cable) or connecting switches or hubs to storage systems
50 m (164 ft)	Within one building, connecting servers to storage systems (adapter must support optical cable) or connecting switches or hubs to storage systems
100 m (328 ft)	Within one complex, connecting servers to storage systems (adapter must support optical cable) or connecting switches or hubs to storage systems
250 m (821 ft,.15 mi)	
500 m (1642 ft,.31 mi)	
Optical cabling is 50 micron (maximum length is 500 m (1,650 ft) or 62.5 micron (maximum length is 300 m (985 ft). Both types are multimode, dual SC, and require a MIA on a DB-9 or hub connector. The minimum length is 2 m (6.8 ft). The minimum bend radius is 3 cm (1.2 in).	

Cable sizes — Copper

Length	Typical use
0.3 m (1 ft), non-equalized	Connecting DPE/DAE and DAE LCCs
1.0 m (3.3 ft), non-equalized	Connecting a hub to an adjacent storage system
3 m (10 ft), non-equalized	Connecting a hub to a storage system in the same cabinet, or daisy chaining from one cabinet to an adjacent cabinet
5 m (16.5 ft), non-equalized	Connecting a hub in one rack to a storage system in another cabinet
10 (33 ft), non-equalized	Connecting servers to hubs and/or storage systems — maximum length for non-equalized copper cable, maximum length between LCCs
30 m (98.5 ft), equalized	Connecting servers to hubs and/or storage systems – maximum length for copper cable
Copper cabling is shielded, 75 ohm twin-axial, shield bonded to DB-9 plug connector shell (360°) FC-AL Standard, Revision 4.4 or higher.	

Component planning diagrams and worksheets follow.

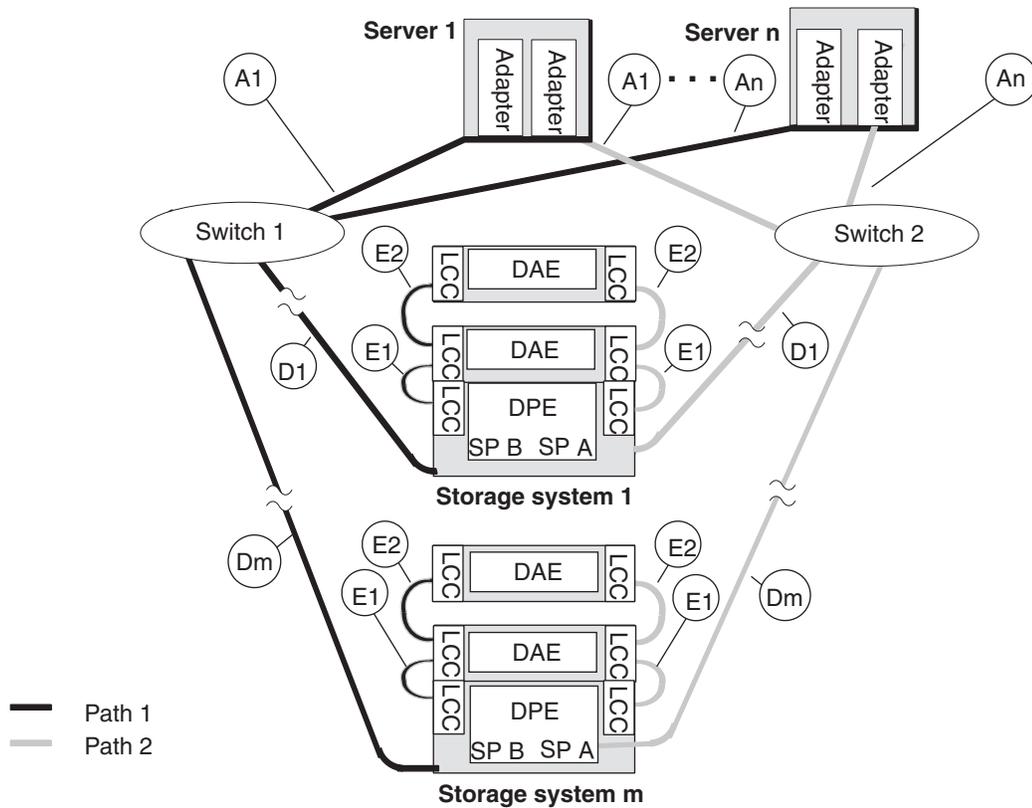
Hardware planning worksheets

Following are worksheets to note the hardware components you want. There are two types of configuration:

- Shared storage
- Unshared storage

Hardware for shared storage

Cable identifier — DPE-based system for shared storage



The cable identifiers used above apply to shared and unshared storage systems. The worksheet applies to shared storage only.

Hardware component worksheet for shared storage

Number of servers:____ Adapters in servers:____ Switches: 16-port:____ 8-port:____ GBICs (1 per port):____

Rackmount DPEs:____ SP/LCC pairs:____ PSs:____ SPSs:____ Rackmount cabinets:____

Rackmount DAEs:____ LCCs:____ PSs:____

Cables between server and switch - Cable A, optical only

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 switches and storage systems - Cable D, copper or optical

Cable D₁, Optical: Number:____Length_____m or ft

Cable D₂, Optical: Number:____Length_____m or ft

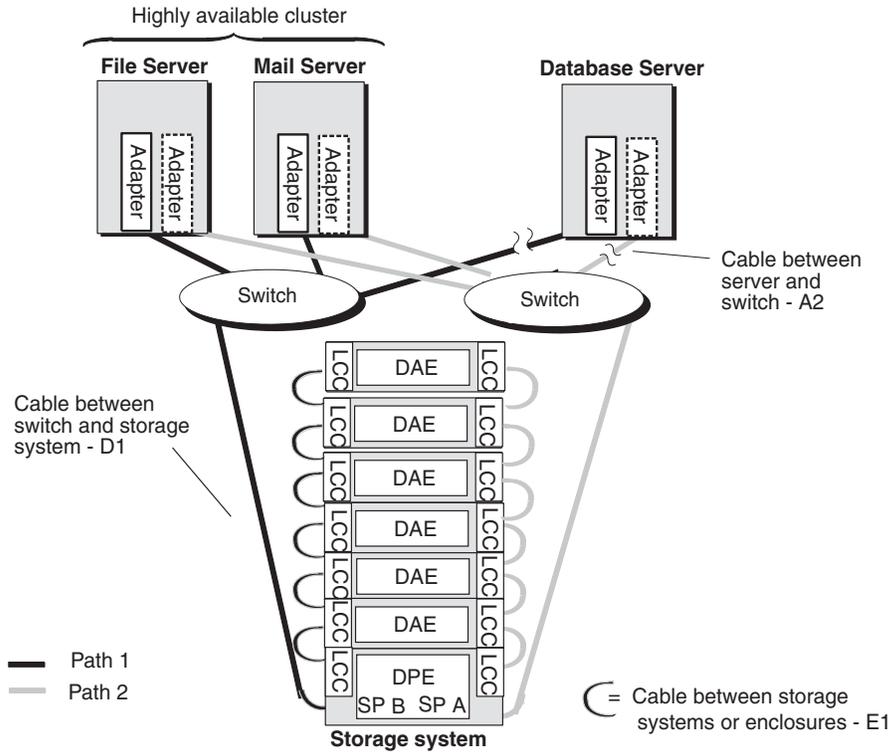
Cable D_m, Optical: Number:____Length_____m or ft

Cables between enclosures - Cable E, which connects LCCs; between a DPE LCC and a DAE LCC, Cable E must be copper; between DAE LCCs, it can be copper or optical.

Cable E₁: Number:____ Copper Optical (for DAE to DAE only).Length_____m or ft

Cable E₂: Number:____ Copper Optical (for DAE to DAE only).Length_____m or ft

Sample shared storage installation and worksheet



Hardware component worksheet for shared storage

Number of servers: 3 Adapters in servers: 6 Switches: 16-port: 8-port: 2 GBICs (1 per port): 6

Rackmount DPEs: 1 SP/LCC pairs: 2 PSs: 2 SPSs: 2 Rackmount cabinets: 1

Rackmount DAEs: 6 LCCs: 12 PSs: 12

Cables between server and switch - Cable A, optical only

Cable A₁, Optical: Number: 2Length 33 m or ft

Cable A₂, Optical: Number: 4Length 1628 m or ft

Cable A_n, Optical: Number: Length m or ft

Cables between switches and storage systems - Cable D, copper or optical

Cable D₁, Optical: Number: 2Length 33 m or ft

Cable D₂, Optical: Number: Length m or ft

Cable D_m, Optical: Number: Length m or ft

Cables between enclosures - Cable E, which connects LCCs; between a DPE LCC and a DAE LCC, Cable E must be copper; between DAE LCCs, it can be copper or optical.

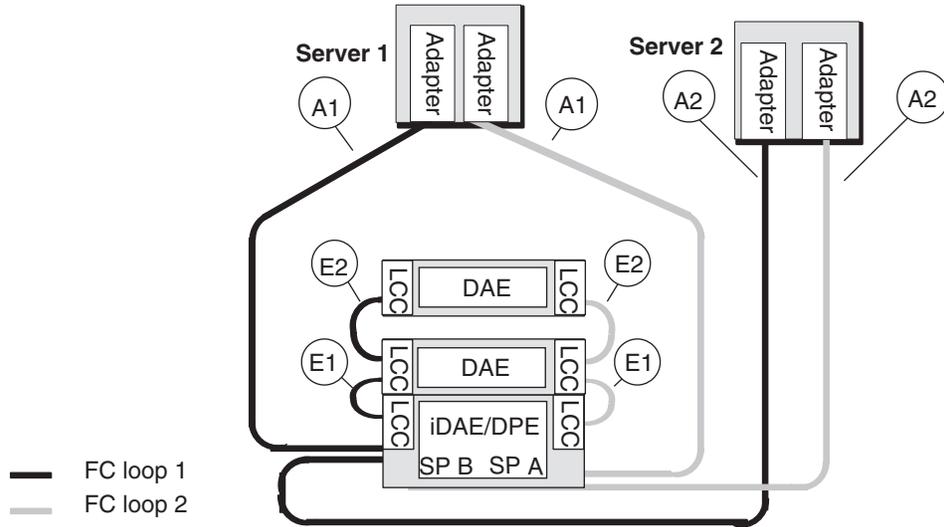
Cable E₁: Number: 12 Copper Optical (for DAE to DAE only)Length 1 m or ft

Cable E₂: Number: Copper Optical (for DAE to DAE only).Length m or ft

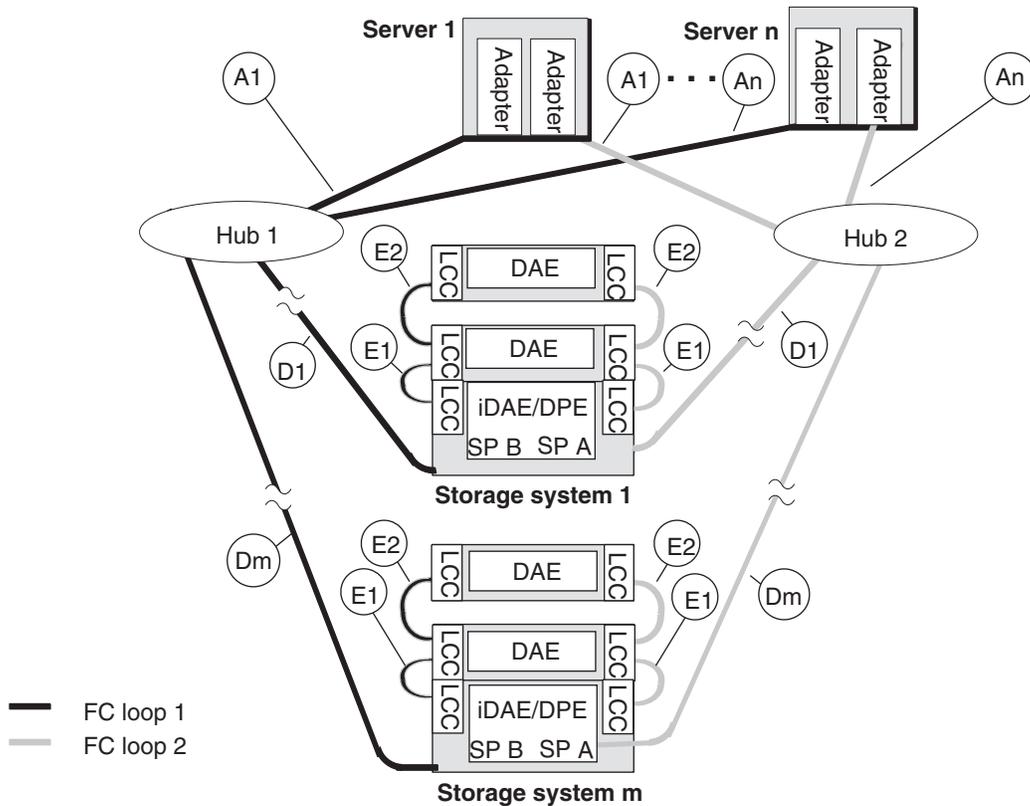
Hardware for unshared storage

The cable identifiers used below and on the following worksheets apply to all types of unshared storage systems. So, if you want to plan a site with different types of systems, you can consolidate all your unshared storage component entries (from the different system types on a single worksheet).

Cable identifier — unshared full-fibre system without hubs



Cable identifier — unshared full-fibre system with hubs



Hardware component worksheet for unshared storage

Number of servers:_____ Adapters in servers:_____ Hubs (copper):_____ MIAs (copper to optical):_____

DPE-based and DAE-only storage systems:

Rackmount DPEs:_____ SP/LCC pairs:_____ PSs:_____ SPSs:_____ Rackmount cabinets:_____

Rackmount iDAEs:_____ SPs:_____ PSs:_____ SPSs:_____ Rackmount cabinets:_____

Rackmount DAEs:_____ LCCs:_____ PSs:_____

Deskside DPEs:_____ SP/ LCC pairs:_____ DAE LCCs:_____ DPE PSs:_____ DAE PSs:_____ SPSs:_____

Deskside iDAEs:_____ SPs:_____ DAE LCCs:_____ PSs:_____ SPSs:_____

Deskside DAEs: 30-slot_____ 10-slot_____ LCCs:_____ PSs:_____

30-slot SCSI-disk storage systems:

Rackmount:_____ SPs:_____ VSCs:_____ BBUs:_____ Rackmount Cabinets:_____

Deskside:_____ SPs:_____ VSCs:_____ BBUs:_____

Cables between server and storage system or between server and hub - Cable A, copper or optical

Cable A₁:Number:_____ Copper Optical Length_____m or ft

Cable A₂:Number:_____ Copper Optical Length_____m or ft

Cable A_n:Number:_____ Copper Optical Length_____m or ft

Cables between hubs and storage systems - Cable D, copper or optical

Cable D₁:Number:_____ Copper Optical..... Length_____m or ft

Cable D₂:Number:_____ Copper Optical..... Length_____m or ft

Cable D_m:Number:_____ Copper Optical Length_____m or ft

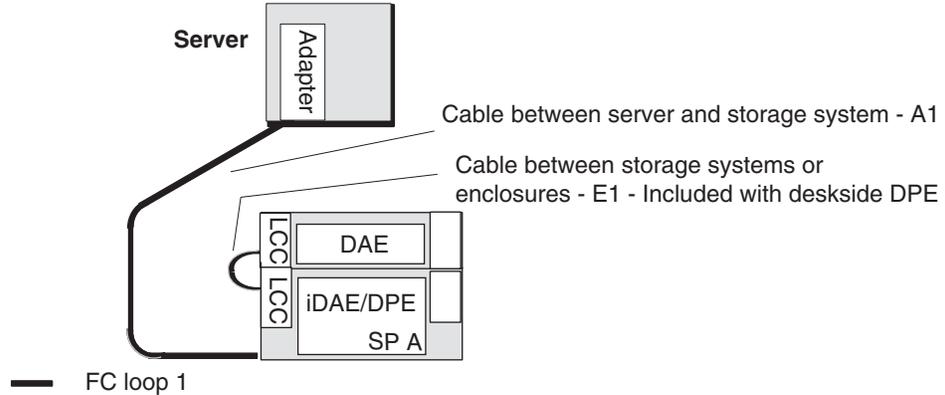
Cables between storage systems or enclosures - Cable E, which connects LCCs or SP-LCC; between a DPE LCC or iDAE SP and a DAE LCC, Cable E must be copper; between DAE LCCs, it can be copper or optical.

Cable E₁:Number:_____ Copper Optical(DAE to DAE only)..... Length_____m or ft

Cable E₂:Number:_____ Copper Optical(DAE to DAE only)..... Length_____m or ft

*Please specify all storage-system components you need, even though you will not need to order them separately, since most or all components will be included with the model of each system you order.

Sample unshared desktide system — basic configuration



Sample component worksheet

Component worksheet for any storage system

Number of servers: 1 Adapters in servers: 1 Hubs (copper): _____ MIAs (copper to optical): _____

DPE-based and DAE-only storage systems:

Rackmount DPEs: _____ SP/LCC pairs: _____ PSs: _____ SPSs: _____ Rackmount cabinets: _____

Rackmount iDAEs: 1 SPs: 1 PSs: 1 SPSs: _____ Rackmount cabinets: 1

Rackmount DAEs: 1 LCCs: 1 PSs: 1

Desktide DPEs: _____ SP/ LCC pairs: _____ DAE LCCs: _____ DPE PSs: _____ DAE PSs: _____ SPSs: _____

Desktide iDAEs: _____ SPs: _____ DAE LCCs: _____ PSs: _____ SPSs: _____

Desktide DAEs: 30-slot _____ 10-slot _____ LCCs: _____ PSs: _____

Cables between server and storage system or between server and hub - Cable A, copper or optical

Cable A₁: Number: 1 Copper Optical Length 10 m or ft

Cable A₂: Number: _____ Copper Optical Length _____ m or ft

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

Cables between hubs - Cable C, copper or optical (maximum 1 between hubs)

Cable C: Number: _____ Copper Optical Length _____ m or ft

Cables between hubs and storage systems - Cable D, copper or optical

Cable D₁: Number: _____ Copper Optical Length _____ m or ft

Cable D₂: Number: _____ Copper Optical Length _____ m or ft

Cable D_m: Number: _____ Copper Optical Length _____ m or ft

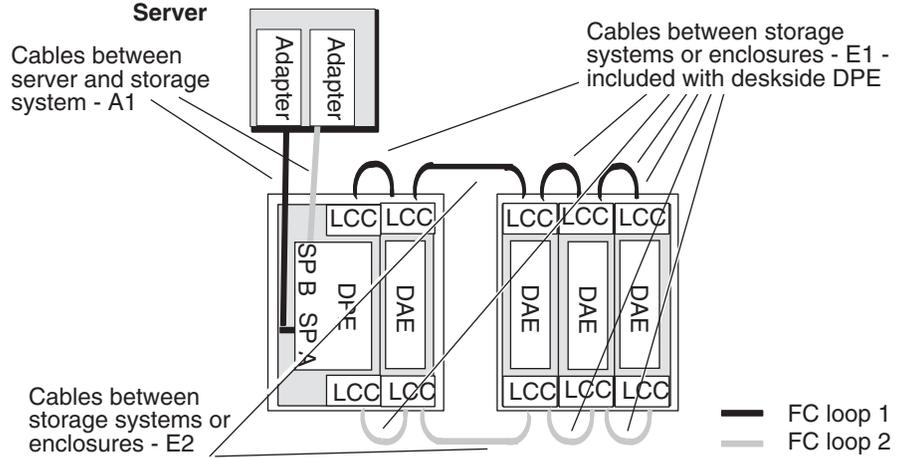
Cables between storage systems or enclosures - Cable E, which connects LCCs or SP-LCC; between a DPE LCC or iDAE SP and a DAE LCC, Cable E must be copper; between DAE LCCs, it can be copper or optical.

Cable E₁: Number: 1 Copper Optical (DAE to DAE only) Length 1 m or ft

Cable E₂: Number: _____ Copper Optical (DAE to DAE only) Length _____ m or ft

*Please specify all storage-system components you need, even though you will not need to order them separately, since most or all components will be included with the model of each system you order.

Sample unshared desktide system — dual-adapter/dual-SP configuration



Sample component worksheet

Hardware component worksheet for unshared storage

Number of servers: 1 Adapters in servers: 2 Hubs (copper): _____ MIAs (copper to optical): _____

DPE-based and DAE-only storage systems:

Rackmount DPEs: _____ SP/LCC pairs: _____ PSs: _____ SPSs: _____ Rackmount cabinets: _____

Rackmount iDAEs: _____ SPs: _____ PSs: _____ SPSs: _____ Rackmount cabinets: _____

Rackmount DAEs: _____ LCCs: _____ PSs: _____

Desktide DPEs: 1 SP/LCC pairs: 2 DAE LCCs: 2 DPE PSs: 2 DAE PSs: 2 SPSs: 1

Desktide iDAEs: _____ SPs: _____ DAE LCCs: _____ PSs: _____ SPSs: _____

Desktide DAEs: 30-slot 1 10-slot _____ LCCs: 6 PSs: 6

Cables between server and storage system or between server and hub - Cable A, copper or optical

Cable A₁: Number: 2 Copper Optical Length 10 m or ft

Cable A₂: Number: _____ Copper Optical Length _____ m or ft

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

Cables between hubs and storage systems - Cable D, copper or optical

Cable D₁: Number: _____ Copper Optical Length _____ m or ft

Cable D₂: Number: _____ Copper Optical Length _____ m or ft

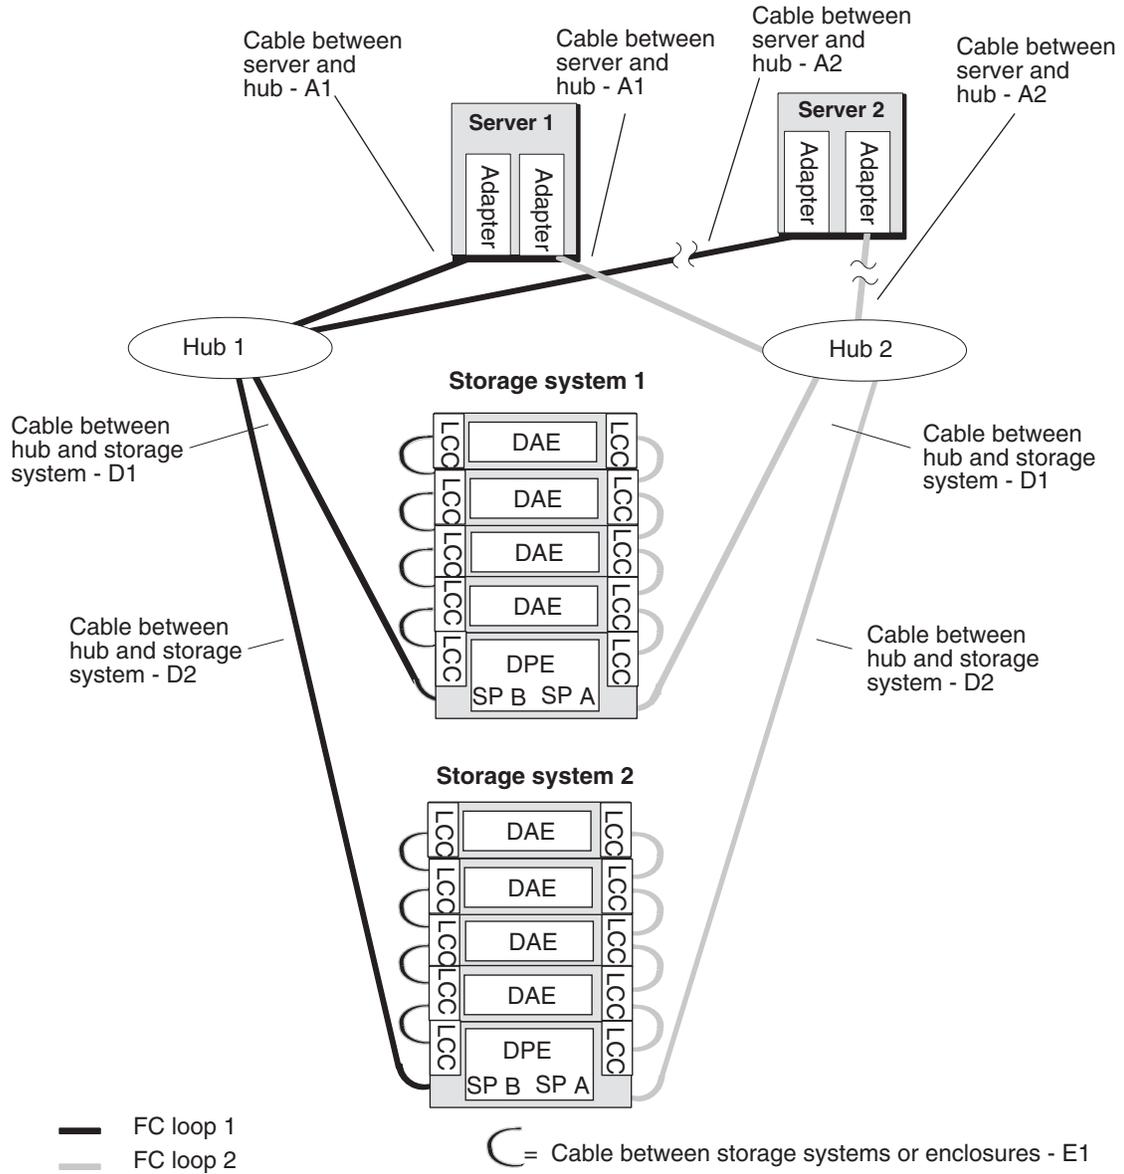
Cable D_m: Number: _____ Copper Optical Length _____ m or ft

Cables between storage systems or enclosures - Cable E, which connects LCCs or SP-LCC; between a DPE LCC or iDAE SP and a DAE LCC, Cable E must be copper; between DAE LCCs, it can be copper or optical.

Cable E₁: Number: 6 Copper Optical (DAE to DAE only) Length 1 m or ft

Cable E₂: Number: 2 Copper Optical (DAE to DAE only) Length 5 m or ft

Sample unshared DPE-based system with hubs — two loops



Sample component worksheet for DPE-based system with hubs — two loops

Component worksheet for any storage system

Number of servers: 2 Adapters in servers: 4 Hubs: 2 MIAs (optical to copper): _____

DPE-based and DAE-only storage systems:

Rackmount DPEs: _____ SP/LCC pairs: 4 PSs: 4 SPSs: 2 Rackmount cabinets: 2
 Rackmount iDAEs: _____ SPSs: _____ PSs: _____ SPSs: _____ Rackmount cabinets: _____
 Rackmount DAEs: 8 LCCs: 16 PSs: 16

Deskside DPEs: _____ SP/ LCC pairs: _____ DAE LCCs: _____ DPE PSs: _____ DAE PSs: _____ SPSs: _____
 Deskside iDAEs: _____ SPSs: _____ DAE LCCs: _____ PSs: _____ SPSs: _____
 Deskside DAEs: 30-slot _____ 10-slot _____ LCCs: _____ PSs: _____

30-slot SCSI-disk storage systems:

Rackmount: _____ SPSs: _____ VSCs: _____ BBU: _____ Rackmount Cabinets: _____
 Deskside: _____ SPSs: _____ VSCs: _____ BBU: _____

Cables between server and storage system or between server and hub - Cable A, copper or optical

Cable A₁: Number: 2 Copper Optical Length 20 m or ft
 Cable A₂: Number: 2 Copper Optical Length 10 m or ft
 Cable A_n: Number: _____ Copper Optical Length _____ m or ft

Cables between hubs and storage systems - Cable D, copper or optical

Cable D₁: Number: 2 Copper Optical Length 20 m or ft
 Cable D₂: Number: 2 Copper Optical Length 20 m or ft
 Cable D_m: Number: _____ Copper Optical Length _____ m or ft

Cables between storage systems or enclosures - Cable E, which connects LCCs or SP-LCC; between a DPE LCC or iDAE SP and a DAE LCC, Cable E must be copper; between DAE LCCs, it can be copper or optical.

Cable E₁: Number: 16 Copper Optical (DAE to DAE only) Length 1 m or ft
 Cable E₂: Number: _____ Copper Optical (DAE to DAE only) Length _____ m or ft

Please specify all storage-system components you need, even though you will not need to order them separately, since most or all components will be included with the model of each system you order.

What next?

This chapter explained hardware components of shared and unshared storage systems. If you have completed the worksheets to your satisfaction, you are ready to consider ordering some of this equipment. Or you may want to read about storage management in the next chapter.



6

Storage-system management

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

- Managing shared or unshared storage systems using Navisphere® Manager software
- Managing unshared storage systems using Navisphere Supervisor software
- Monitoring DAE-only storage systems (JBODs)
- Storage management worksheets

Navisphere software lets you bind and unbind disks, 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.

Navisphere products have two parts: a graphical user interface (GUI) and an Agent. The GUIs run on a management station, accessible from a common framework, and communicate with storage systems through a single Agent application that runs on each server. The management station and server are often separate systems, but they need not be; that is, you can run one or more GUIs and an agent on one server. The Navisphere products are

- Navisphere Manager, which lets you manage multiple storage systems on multiple servers simultaneously; Navisphere Manager is required for shared storage and optional for unshared storage.
- Navisphere Supervisor software, which lets you manage one storage system on one server at a time, which is included with each storage system. Navisphere Supervisor supports unshared storage systems only.
- Navisphere Analyzer, which lets you measure, compare, and chart the performance of SPs, LUNs, and disks.
- Navisphere Integrator, which provides an interface between Navisphere products and HP OpenView and Tivoli.
- Navisphere Event Monitor, which checks storage systems for fault conditions and can notify you and/or customer service if any fault condition occurs.
- Navisphere Agent, which is included with each storage system, and Navisphere CLI (Command Line Interpreter), which lets you bypass the GUIs and issue commands directly to storage systems.

Managing shared or unshared storage systems using Navisphere Manager software

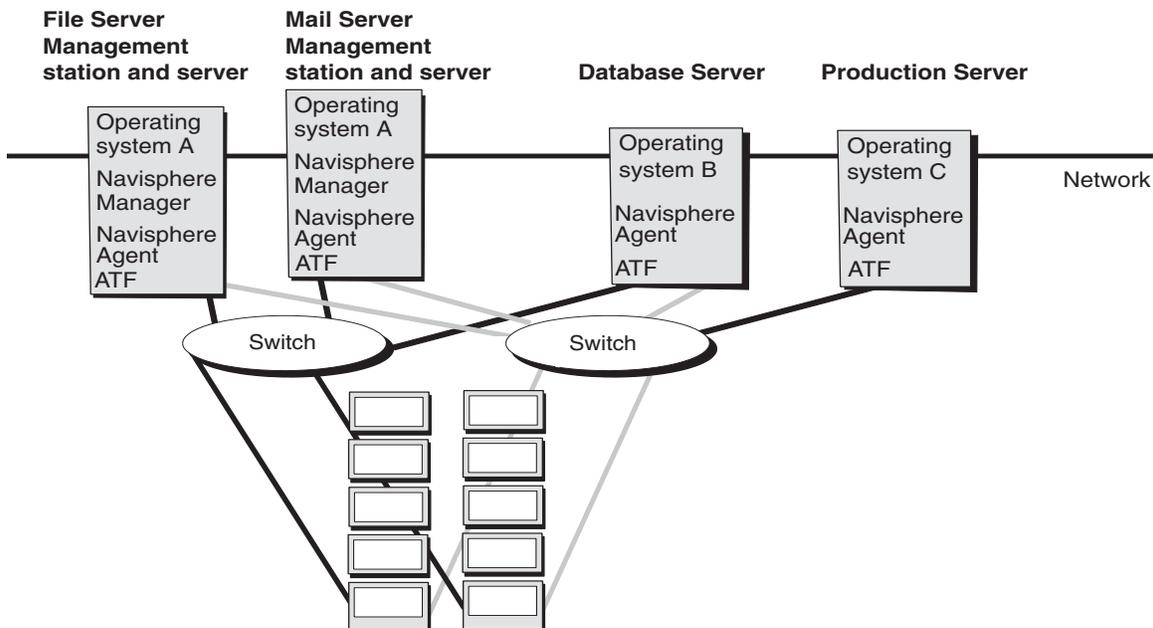
Navisphere Manager software (Manager) lets you manage multiple storage systems connected to servers on a TCP/IP network. Manager offers extensive management features and includes comprehensive on-line help. Manager is required for shared storage and optional for unshared storage.

Manager runs on a management station which is a Windows NT® host. The servers connected to a storage system can run Windows NT or one of several UNIX® operating systems. With shared storage, servers connected to the SAN can run different operating systems; with unshared storage, servers connected to the same storage system must run the same operating system.

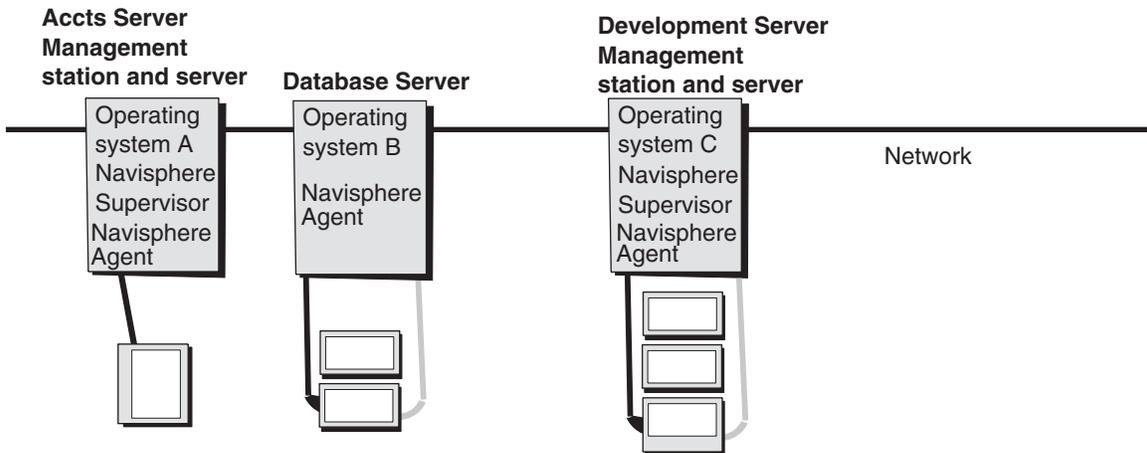
Another software component for shared storage or other high-availability storage systems is Application-Transparent Failover (ATF). ATF is a software product that works with shared storage or other high-availability storage systems to let applications continue running after the failure in the path to a LUN. If a host-bus adapter, switch, hub, SP, or cable fails, ATF can automatically route I/O to the LUNs the applications need.

The following figures show Navisphere Manager in shared and unshared environments.

Sample shared environment (SAN) with Navisphere Manager



Sample unshared environment with Navisphere Manager

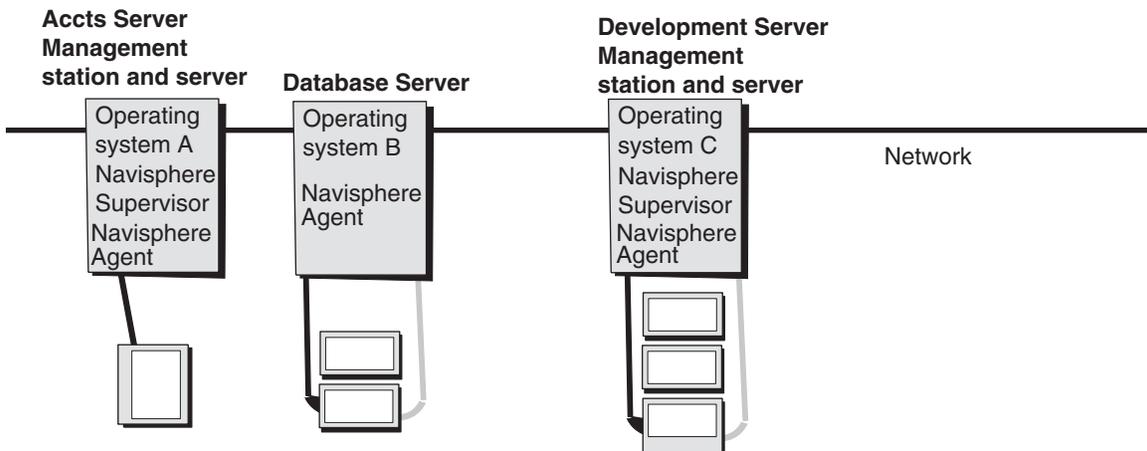


Managing unshared storage systems using Navisphere Supervisor software

Navisphere Supervisor software, like Manager, consists of a graphical user interface, called Supervisor, and a storage-system agent, called the Navisphere Agent — the same Agent used by Manager. However, Supervisor software lets you manage only one storage system at a time.

The sample management environment below shows a Windows NT server that is both a manager and an agent. You can use this server to manage its own storage system and/or to manage the storage systems on the other server with agents.

Sample unshared environment with Navisphere Supervisor



Monitoring DAE-only storage systems (JBODs)

With a DAE-only (JBOD) system, you can monitor status with Navisphere Manager or Supervisor. To create LUNs of different RAID types, you must use the software RAID utility designed for the operating system.

Storage management worksheets

This section includes two worksheets: one for shared storage and one for unshared storage. The following worksheet will help you plan your storage system management environment. For each host, complete a section.

For the shared storage worksheet, complete the management station hostname and operating system; then decide whether you want the Navisphere Analyzer and/or Event Monitor and, if so, mark the appropriate boxes. Then write the name of each managed server, with operating system, storage group, and configuration access specification. You can copy much of the needed information from the LUN and Storage Group planning worksheet in Chapter 3.

Management Utility Worksheet – Shared Storage

Management station hostname: _____ Operating system: _____

Software: Navisphere Manager/Navisphere Agent Navisphere Analyzer Navisphere Event Monitor

List all the servers this host will manage. Each managed server must run an Agent and ATF software of the same type as its operating system.

Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access

Management station hostname: _____ Operating system: _____

Software: Navisphere Manager/Navisphere Agent Navisphere Analyzer Navisphere Event Monitor

List all the servers this host will manage. Each managed server must run an Agent and ATF software of the same type as its operating system.

Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access
Server: _____ Op sys: _____	Storage Group number or name: _____ <input type="checkbox"/> Config Access

For unshared storage, for each host, choose a Navisphere product. The host may be a management station that is not a server (complete only the Manager section); it may be a management station that is a server (complete the Manager section and mark the Agent box), or it may be a server (mark the Agent box).

Management Utility Worksheet – Unshared Storage

Hostname: _____ Operating system: _____

Storage system type: DPE-based iDAE-based 30-slot SCSI disk

Software: Navisphere Manager and Navisphere Agent Navisphere Supervisor and Navisphere Agent

List all the servers this host will manage. Each managed server must run an Agent of the same type as its operating system.

Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____

Hostname: _____ Operating system: _____

Storage system type: DPE-based iDAE-based 30-slot SCSI disk

Software: Navisphere Manager and Navisphere Agent Navisphere Supervisor and Navisphere Agent

List all the servers this host will manage. Each managed server must run an Agent of the same type as its operating system.

Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____

Hostname: _____ Operating system: _____

Storage system type: DPE-based iDAE-based 30-slot SCSI disk

Software: Navisphere Manager and Navisphere Agent Navisphere Supervisor and Navisphere Agent

List all the servers this host will manage. Each managed server must run an Agent of the same type as its operating system.

Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____
Server: _____ Oper sys: _____	Server: _____ Oper sys: _____



Index

Within this index, a range of page numbers indicates that the reference spans those pages.

Numerics

- 30-slot SCSI-disk storage systems
 - dimensions 5-15
 - disk structure example
 - dual-server 4-6
 - single-server 4-5
 - physical disk structure example
 - dual-server 4-7
 - site requirements 5-15
 - weight 5-15

A

- ac power requirements
 - 30-slot SCSI-disk systems 5-15
 - DAE-only storage system 5-14
 - DPE storage system 5-12
 - iDAE storage system 5-13
- application planning
 - shared storage 3-5
 - unshared storage 4-8
- application worksheet, completing
 - shared 3-4, 4-8
 - unshared 4-9
- applications, sample. for RAID
 - groups 2-14
- application-transparent failover (ATF)
 - introduced 6-2
- array, *see* disk-array storage system
- attach kit, *see* interface kit
- audience for manual v

B

- basic configuration
 - components 4-2
 - features 4-2

C

- cabinets for rackmount storage
 - systems 5-16
- cabling
 - guidelines 5-16
 - introduced 1-3
 - types and sizes 5-18
- cache
 - about 5-8
 - page size 4-16
- cascading hubs 1-4, 1-6

- communication with storage system,
 - Chapter 6
- configurations
 - error recovery tradeoffs, *see* error recovery features
 - LUN and file system
 - planning 3-4
 - RAID
 - compared 2-9
 - examples 4-5
 - guidelines 2-13–2-14
 - planning 3-4
 - server and storage system
 - unshared 4-1
 - shared storage
 - examples 3-1
 - storage system
 - basic 4-2
 - dual-adapter/dual-SP 4-2
 - unshared 4-1
 - tradeoffs
 - shared storage systems 5-8
 - unshared storage 5-9
 - unshared storage
 - examples 4-1
 - cooling requirements, cabinet 5-16
 - copper cable, types and sizes 5-18
 - CRUs (customer-replaceable units)
 - locating 5-6

D

- DAE-only storage systems
 - dimensions 5-14
 - introduced 1-11, 5-4
 - site requirements 5-14
 - weight 5-14
- data sheets, hardware 5-11
- device name, operating system 3-13, 4-18
- disk
 - configuration types
 - compared 2-9
 - examples 4-5
 - guidelines 2-13–2-14
 - planning 3-4
 - configuration, *see* RAID group
 - mirror, defined 2-2
 - shared storage
 - examples 3-1
 - striping, defined 2-2

- disk (continued)
 - unshared storage examples 4-1
 - Disk Array Enclosure (DAE)
 - site requirements 5-14
 - disk module(s)
 - capacity, defined 4-17–4-18
 - IDs
 - 30-slot SCSI-disk system 5-7
 - full-fibre system 3-2, 4-5
 - disk unit number on worksheet 3-6, 3-11, 4-10, 4-17–4-18
 - disk-array storage system
 - communicating with, Chapter 6
 - managing 6-2
 - managing, Chapter 6
 - types
 - defined 1-10
 - DPE storage systems
 - components 5-2
 - dimensions 5-12
 - site requirements 5-12
 - weight 5-12
 - dual paths to LUNs 3-1, 4-1
 - dual-adapter/dual-SP configuration
 - components 4-2
 - disk structure example 4-5
 - features 4-2
 - dual-initiator/dual-loop configuration
 - disk structure example 4-6–4-7
 - dual-ported disk
 - as function of dual SPs 5-3, 5-7
 - dual-server configurations
 - disk structure example 4-6–4-7
 - unshared storage 4-2
- E**
- enclosure address (EA)
 - DPE 5-2, 5-6
 - error recovery features
 - configurations with hubs 4-4
 - server configurations 4-2
- F**
- Fibre Channel
 - adapter 1-3
 - components 1-3
 - defined 1-1
 - hub, description 1-6
 - switch, description 1-4
 - file system
 - name 3-13, 4-18
 - worksheet
 - completing 3-11, 4-16
 - footprint
 - 30-slot SCSI-disk systems 5-15
 - cabinet 5-16
 - DAE-only storage system 5-14
 - DPE storage systems 5-12
 - iDAE storage systems 5-13
- G**
- GBIC (Gigabit Interface Converter),
 - about 1-4
 - global spare, *see* hot spare
 - grounding requirements 5-16
 - GUI (in storage-system management utilities) 6-2
- H**
- hardware
 - data sheets 5-11
 - mirroring 2-2
 - planning worksheets
 - shared storage 5-19
 - shared storage 5-1–5-2
 - heat dissipation
 - 30-slot SCSI-disk systems 5-15
 - DAE-only storage system 5-14
 - DPE storage system 5-12
 - iDAE storage system 5-13
 - height
 - 30-slot SCSI-disk systems 5-15
 - DAE-only storage systems 5-14
 - DPE storage system 5-12
 - iDAE storage system 5-13
 - high availability options
 - unshared storage system 5-10
 - high-availability features, *see* error recovery features
 - host, *see* server
 - host-bus adapter (HBA)
 - description 1-3
 - hot spare
 - defined 2-8
 - sample applications 2-15
 - when to use 2-14
 - hub
 - configurations 4-4
 - description 1-6
 - planning system with 5-22
 - sample hardware worksheet 5-26
- I**
- iDAE storage systems
 - dimensions 5-13
 - site requirements 5-13
 - weight 5-13
 - image, disk, defined 2-2
 - individual access array, *see* RAID 5 group

individual disk unit
 defined 2-7
 disk space usage 2-13
 performance 2-11
 sample applications 2-15
 when to use 2-14

interconnect components 5-1
 cables, hubs, switches 1-3

interface kit 1-3

L

LCC (link control card)DPE storage system 5-6

logical volume, *see* file system

LUN (logical unit)
 configurations
 compared 2-9
 examples 4-5
 guidelines 2-13–2-14
 individual disk
 defined 2-7
 planning 3-4
 RAID 0
 defined 2-6
 RAID 1 mirrored pair 2-5
 RAID 1/0 group
 defined 2-6
 RAID 3 group
 defined 2-4
 RAID 5 group
 defined 2-3
 shared storage
 examples 3-1
 unshared storage
 examples 4-1
 worksheets 3-9–3-11, 4-14–4-16

disk mirror, defined 2-2

hot spare, defined 2-8

in RAID Group 2-2

number on worksheet 3-6, 3-11, 4-10, 4-17–4-18

SP control of 5-3, 5-8

LUNs, paths to 3-1, 4-1

M

Manager utility 6-1

manual, about v

MIA (media interface adapter), about 1-3

mirrored pair, *see* RAID 1 mirrored pair

mirrored RAID 0 group, *see* RAID 1/0 group

mirroring, defined 2-2

N

Navisphere Manager utility 6-1

Navisphere product set 6-2

node, defined 1-2

nonredundant array, *see* RAID 0 group

O

operating system
 device name for disk unit 3-13, 4-18
 software mirroring 2-2

optical cable, types and sizes 5-18

organization of manual v

P

page size, cache 3-11, 4-16

parallel access array, *see* RAID 3 group

paths to LUNs 3-1, 4-1

performance, RAID group 2-10

physical disk unit, *see* LUN (logical unit)

physical volume, *see* LUN (logical unit)

planning, LUNs and file systems 3-4

plant requirements
 30-slot SCSI-disk systems 5-15
 DAE 5-14
 iDAE 5-13

plug types 5-16

power requirements
 30-slot SCSI-disk systems 5-15
 DAE-only storage system 5-14
 DPEstorage system 5-12
 iDAE storage system 5-13

power supplies (PSs)
 DPEstorage system 5-6

R

rackmount model, DPE storage system 5-2

RAID group
 configurations
 compared 2-9
 examples 4-5
 guidelines 2-13–2-14
 planning 3-4
 performance 2-10
 RAID 3 versus RAID 5 2-5
 shared storage examples 3-1
 types and tradeoffs, 2-1; *see also* Chapter 2
 unshared storage examples 4-1

RAID groups
 types and tradeoffs, 2-1; *see also* Chapter 2

RAID groups and LUNs 2-2

- RAID 0 group
 - defined 2-6
 - sample applications 2-15
 - when to use 2-14
 - RAID 1 mirrored pair
 - defined 2-5
 - sample applications 2-15
 - when to use 2-13
 - RAID 1/0 group
 - defined 2-6
 - sample applications 2-15
 - when to use 2-14
 - RAID 3 group
 - defined 2-4
 - sample applications 2-15
 - when to use 2-13
 - RAID 5 group
 - defined 2-3
 - sample applications 2-14
 - when to use 2-13
 - redundant array of independent disks (RAID), *see* RAID group
- S**
- server
 - cabling guidelines 5-16
 - component 5-1
 - configuration with hubs 4-4
 - configurations
 - unshared storage 4-1
 - connection to storage system, *see* cabling
 - dual-server configuration
 - unshared storage 4-2
 - planning worksheet
 - unshared storage 5-23
 - unshared configurations
 - tradeoffs 4-2
 - service clearance
 - DPE storage system 5-12
 - iDAE storage system 5-13
 - shared storage systems
 - cabinets 5-16
 - components 5-1
 - disk structure example 3-1
 - hardware
 - planning worksheets 5-19
 - single-server configurations
 - disk structure example 4-5
 - features 4-2
 - unshared storage 4-1
 - site requirements
 - 30-slot SCSI-disk systems 5-15
 - DAE 5-14
 - DAE-only storage systems 5-14
 - DPE storage systems 5-12
 - iDAE storage systems 5-13
 - size
 - cache 3-11, 4-16
 - software mirroring 2-2
 - SP (storage processor)
 - description 5-7
 - DPE storage system 5-3
 - FC-AL address ID 3-11, 4-16
 - SPS (standby power supply), DPE storage system 5-6
 - sstorage components 5-1
 - shared storage 5-2
 - storage management worksheets 6-4
 - storage system
 - caching
 - on worksheet 4-18
 - configurations
 - basic 4-2
 - dual-adapter/dual-SP 4-2
 - storage-system caching, as feature 5-8
 - stripe
 - with RAID 1/0, RAID 0 2-7
 - with RAID 5, RAID 3 2-3
 - stripe, defined 2-2
 - Supervisor utility 6-3
 - switch
 - description 1-4
 - in sample shared storage
 - configuration 3-1
 - planning system with 5-19
- T**
- temperature requirements
 - 30-slot SCSI-disk systems 5-15
 - DAE-only storage system 5-14
 - DPE storage system 5-12
 - iDAE storage system 5-13
 - terms
 - RAID 2-1
 - tradeoffs
 - configuration
 - shared storage 5-8
 - shared storage systems 5-9
- U**
- unshared storage
 - disk structure example 4-1
 - hardware
 - planning 5-22
 - storage system configurations 4-1
- V**
- vault disks 2-8
 - volume name 3-13

W

weight

- 30-slot SCSI-disk systems 5-15
- DAE-only storage systems 5-14
- DPE storage system 5-12
- iDAE storage system 5-13
- storage system installation 5-16

worksheet

- application 3-4
- completing 3-4
- component planning
 - shared storage 5-19
 - unshared storage 5-22
- LUN configuration 3-9–3-11, 4-14–4-16
- storage management 6-4

Vos remarques sur ce document / Technical publication remark form

Titre / Title : Bull DAS – Disk Array Storage Systems Planning a DAS Installation Fibre Channel Environments

N° Référence / Reference N° : 86 A1 94JX 04

Daté / Dated : July 2000

ERREURS DETECTEES / ERRORS IN PUBLICATION

AMELIORATIONS SUGGEREES / SUGGESTIONS FOR IMPROVEMENT TO PUBLICATION

Vos remarques et suggestions seront examinées attentivement.

Si vous désirez une réponse écrite, veuillez indiquer ci-après votre adresse postale complète.

Your comments will be promptly investigated by qualified technical personnel and action will be taken as required.

If you require a written reply, please furnish your complete mailing address below.

NOM / NAME : _____ Date : _____

SOCIETE / COMPANY : _____

ADRESSE / ADDRESS : _____

Remettez cet imprimé à un responsable BULL ou envoyez-le directement à :

Please give this technical publication remark form to your BULL representative or mail to:

**BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE**

Technical Publications Ordering Form

Bon de Commande de Documents Techniques

To order additional publications, please fill up a copy of this form and send it via mail to:

Pour commander des documents techniques, remplissez une copie de ce formulaire et envoyez-la à :

BULL CEDOC
ATTN / MME DUMOULIN
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

Managers / Gestionnaires :
Mrs. / Mme : C. DUMOULIN +33 (0) 2 41 73 76 65
Mr. / M : L. CHERUBIN +33 (0) 2 41 73 63 96
FAX : +33 (0) 2 41 73 60 19
E-Mail / Courrier Electronique : srv.Cedoc@franp.bull.fr

Or visit our web site at: / Ou visitez notre site web à:

<http://www-frec.bull.com> (PUBLICATIONS, Technical Literature, Ordering Form)

CEDOC Reference # N° Référence CEDOC	Qty Qté	CEDOC Reference # N° Référence CEDOC	Qty Qté	CEDOC Reference # N° Référence CEDOC	Qty Qté
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
__ __ __ __ __ [__]		__ __ __ __ __ [__]		__ __ __ __ __ [__]	
[__]: no revision number means latest revision / pas de numéro de révision signifie révision la plus récente					

NOM / NAME : _____ Date : _____

SOCIETE / COMPANY : _____

ADRESSE / ADDRESS : _____

PHONE / TELEPHONE : _____ FAX : _____

E-MAIL : _____

For Bull Subsidiaries / Pour les Filiales Bull :

Identification: _____

For Bull Affiliated Customers / Pour les Clients Affiliés Bull :

Customer Code / Code Client : _____

For Bull Internal Customers / Pour les Clients Internes Bull :

Budgetary Section / Section Budgétaire : _____

For Others / Pour les Autres :

Please ask your Bull representative. / Merci de demander à votre contact Bull.

BULL CEDOC
357 AVENUE PATTON
B.P.20845
49008 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A1 94JX 04

PLACE BAR CODE IN LOWER
LEFT CORNER



Utiliser les marques de découpe pour obtenir les étiquettes.
Use the cut marks to get the labels.

┌ ───┐ ┌ ───┐
DAS
Planning a DAS
Installation
Fibre Channel
Environments
86 A1 94JX 04
└ ───┘ └ ───┘

┌ ───┐ ┌ ───┐
DAS
Planning a DAS
Installation
Fibre Channel
Environments
86 A1 94JX 04
└ ───┘ └ ───┘

┌ ───┐ ┌ ───┐
DAS
Planning a DAS
Installation
Fibre Channel
Environments
86 A1 94JX 04
└ ───┘ └ ───┘