014–2373–04

DAS – Disk-Array Storage System

Planning a DAS Installation (SCSI Environments)

ORDER REFERENCE
86 A1 84GX 01

DAS – Disk-Array Storage System

Planning a DAS Installation (SCSI Environments)

June 1998

BULL ELECTRONICS EUROPE S.A.
Service CEDOC
331 Avenue PATTON – BP 428
49004 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A1 84GX 01

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

This book is derived from the book "Disk-Array Storage Systems Configuration Planning – SCSI Environments", Data General Corporation.

Contents

About DAS disk-array storage systems

A DAS disk-array storage system provides enormous disk storage capacity and highly available data at low cost. DAS storage systems are available in a deskside or cabinet-mountable chassis, with slots to hold 30, 20, 10, or 7 disks (the 7-slot model is available in a deskside chassis only). They offer ideal disk storage for servers and personal computers.

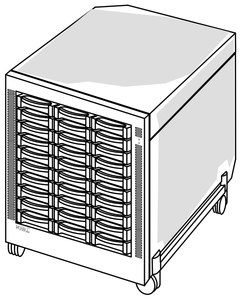

DAS systems let you configure terabytes of disk storage on high-end server systems. And you can easily make this storage highly available using disk arrays and clusters of servers.

Please read this guide

- if you are considering purchase of a DAS disk–array storage system and want to understand the versatility and expandability they offer; or
- before you install a DAS disk-array storage system.

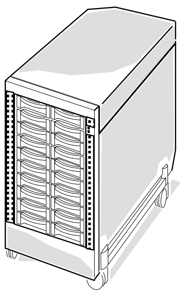

You do not need to own a DAS storage system to use the guide. But you should be familiar with the server(s) that will host the storage system and with the operating system the server(s) will run.

After reading this guide, you will be able to decide which storage system and disk configuration you want, and plan them. The guide describes

- Storage-system features and benefits
- Storage-system components
- Communication with a storage system
- Server and storage-system configurations
- Disk configuration definitions
- Disk configuration benefits and tradeoffs
- Site requirements
- Configuration planning worksheets

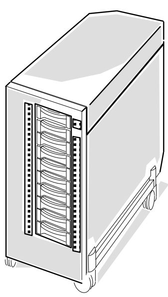

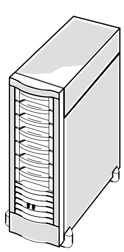

30-slot disk-array storage system (3x00 series)

Deskside model

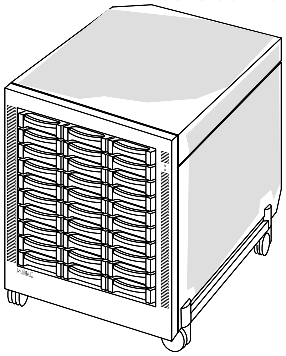

Rackmount model

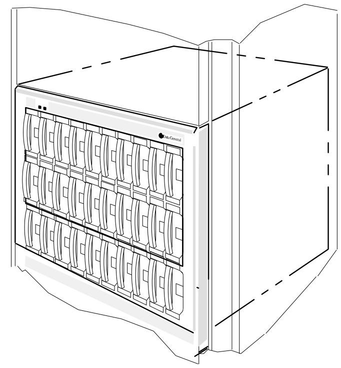

20-slot disk-array storage system (2x00 series)

Deskside model

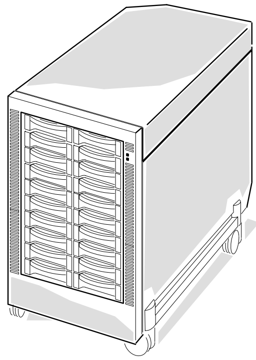

Rackmount model

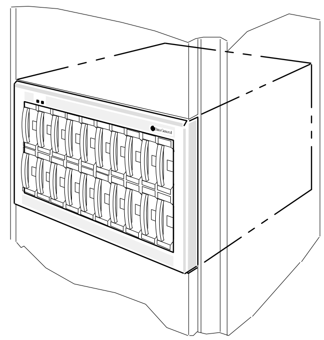

10-slot disk-array storage system (1x00 series)

Deskside model

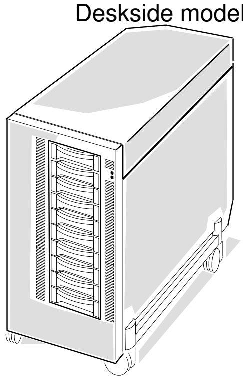

Rackmount model

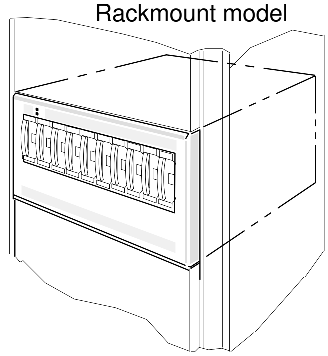

7-slot disk-array storage system (150 series)

Deskside model

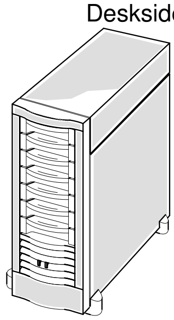

Storage-system features and benefits

The disk-array storage-system features and benefits are

- High-capacity disk storage
- Highly available data
- Storage-system caching
- Low cost of ownership

High-capacity disk storage

A storage system provides lots of disk storage for many users and their applications. With 9-Gbyte disk modules (8.8 formatted), a 30-slot storage system can store 264 Gbytes of data or 211 Gbytes in highly available form; a 20-slot system can store 176 Gbytes traditional or 141 Gbytes in highly available form. Three 30-slot or four 20-slot deskside systems fit in a cabinet, providing 634 or 563 Gbytes, respectively, of highly available storage in a footprint of less than five square feet (.55 square meters).

The disk-array storage systems use Redundant Array of Independent Disks (RAID) technology to group disks into arrays that appear to the operating system as single disk units. A RAID group can include as many as 16 disks, which, with 9-Gbyte disks, can hold a highly available database of 132 Gbytes. This lets you tailor disk storage space to your applications' specific needs, and provide room for growth without acquiring expensive excess space.

Highly available data

Highly available data let applications continue running after hardware failure. The storage system makes data highly available using RAID groups, automatic transfer of disk units on hardware failure, and redundant components that you can replace under power.

RAID groups for data redundancy

RAID technology helps storage systems provide highly available data. RAID uses parity data to let a disk array continue without data loss after a disk failure. It uses the parity information on the working disks to maintain an accurate copy of the array's data. After the failed disk is replaced, the storage system rebuilds the array from the information stored on other disks. The rebuilt array contains an exact replica of the information it would have held if the module did not fail.

You can extend the availability of all RAID groups in a storage system by using one or more disk modules as hot spares. During normal operation, a hot spare stores no data but offers additional redundancy: the storage system will integrate the hot spare automatically if any module in a RAID group fails. A hot spare is global, so the disk-array storage system will integrate it in any RAID group that suffers a failure.

Transfer of disk units on hardware failure (failover)

A storage system can connect to one or two servers in different ways, depending on the level of availability you need. In the basic configuration, one server has one SCSI-2 adapter and one storage processor (SP). For higher availability, a server can use two adapters and two SPs. With this backup hardware, if an adapter or processor fails, all I/O traffic can be rerouted over the surviving I/O path, thus maintaining access to all the storage-system disk units.

For even higher availability, two servers can connect to one storage system. This provides a backup server if one server fails. As with a single server, I/O can be rerouted over a working I/O path (to another cluster server), maintaining access to all the storage-system disk units.

Another benefit with two servers is the option of using a failover software, such as HACMP for AIX. A program on each server monitors the other server over a network; if one server does not respond, the program in the other server runs scripts that take over the disks and restart applications. The transfer and restart occur automatically, without human intervention.

Another failover option is application-transparent failover (ATF). ATF software lets applications continue running without human intervention after SCSI-2 adapter, SCSI-bus, or SP failure.

Redundant hardware that you can replace under power

Storage-system redundancy extends beyond RAID to hardware other than disk units. It can support a second SP (30-, 20-, and 10-slot models only) and redundant power supply. With a redundant component, the storage system can continue after failure of one component. These customer-replaceable units (CRUs) include disk modules, SPs, power supplies (VSCs), fan module, and battery backup unit (BBU).

Storage-system caching

Storage-system read and write caching improve I/O performance for several types of RAID groups. Write caching, particularly, helps write performance — an inherent problem for RAID types that require writing to multiple disk modules. Read and write caching enhance performance in two ways:

- For a read request — If the request seeks information that's already in the read or write cache, the storage-system can deliver the information immediately – much faster than a disk access can. Read caching offers the following added features and benefits:
  - a prefetch capability that lets the SP anticipate the data you will need and read it from disk into the read cache; this helps performance significantly with sequential I/O operations. You can set prefetch parameters to match demands of custom applications.
  - independent cache space for read I/O; this eliminates the need to mirror read data in the cache to another SP. With two SPs, each SP can maintain its own independent read cache.
- For a write request — The storage system writes information to write cache memory instead of disk, allowing the server to continue processing as if the disk write had completed. The disk write from cache memory occurs later, at the most expedient time. If the request modifies information that's already in cache memory waiting to be written to disk, the storage system updates the information in the cache before writing it; this requires just one disk write instead of two.

Read caching requires at least 8 Mbytes of SIMM memory on at least one SP. For write caching, any storage system except a 7-slot requires two SPs, a battery-backup unit (BBU), and disk modules in five specific slots. For write caching, a 7-slot system requires an SP with a battery-equipped PCMCIA memory card to maintain a backup image of the cache information.

You can allot the SP SIMM memory to read caching, write caching, or both – after subtracting a fixed amount for storage-system use.

Parallel access disk arrays (known as RAID-3 groups), do not use storage-system caching; instead, RAID-3 groups use a different buffering mechanism for high performance.

Write cache integrity

In a 30-, 20,- or 10- slot storage system, data in the SP write cache is protected from power loss by a battery-backup unit (BBU). If line power fails, BBU power lets the storage system write the contents of the write cache to the vault disks. Vault disks are standard disk modules that have reserved space outside operating system control. When power returns, the storage system automatically reads data from the vault disks into the write cache, and then writes it to the file systems whose data was cached. This design ensures that all write cached information reaches its destination.

Vault disk use is independent of user data storage; a module's role as a vault disk has no effect on its storage capacity or performance.

Storage-system write caching works with two SPs or one SP. For high availability with 30-, 20-, or 10- slot systems, you must use two SPs. With two SPs, the SPs mirror the write cache data, so if one SP fails or encounters memory errors, the other SP maintains write cache integrity.

> **IMPORTANT** Without a second SP to mirror the cache, an SP failure causes loss of the write cache contents. If the write cache contents are lost, all disks to which write caching is enabled lose the integrity of their file systems. A system operator must recreate each file system and reload its information from the latest backup. To avoid this downtime, we recommend using two SPs for write caching with 30-, 20-, and 10-slot systems.

Write cache integrity in 7-slot systems

In a 7-slot storage system, the SP mirrors its write cache on a PCMCIA memory card that has its own replaceable backup battery. If line power or the SP fails, the PCMCIA card retains the cache information. When power returns, or you insert the PCMCIA card in a replacement SP, the storage system automatically writes the cached data from the card to the disks. This mechanism ensures that all write cached data reaches its destination.

Low cost of ownership

The storage system keeps your ownership costs low. You can bind arrays of different sizes and types for different needs. For example, you can create a two-disk RAID array (mirror) of 4-Gbyte disk modules for one file system and a five-disk RAID array of 9-Gbyte modules for another file system.

As needs change, you can expand a storage system for greater capacity and/or higher availability. You can start with an entry level storage system and three disk modules in a basic configuration, and then upgrade in steps to a full complement of disk modules (30, 20, 10, or 7, depending on model). With 30-, 20-, or 10-slot models, you can add a second SCSI-2 adapter and SP, caching, redundant power supply, and automatic transfer on failure. With 7-slot models, you can add SP memory and an PCMCIA memory card to provide write caching. All these enhancements (except for the adapter and configuration change) can occur while the storage system is under power, running applications.

Replacing a failed component — disk, storage processor, power supply, fan, or battery backup unit — is easy. A person who does not understand computer hardware can replace the failed component at his or her convenience, while the storage system is running and serving users. Replacing a component requires no mechanical expertise or tools, except a screwdriver for external cabling.

These disk-array storage systems are compact, with modest power requirements. The largest rackmount storage system, the 30-slot model, measures 47 x 76 x 36 cm (18.4 x 30 x 14 inches); the largest deskside storage systems, with 20 slots, measures 63 x 76 x 36 cm (24.8 x 30 x 14 inches). Three 30-slot or four 20-slot deskside systems fit in a cabinet, providing — with 9-Gbyte disks (8.8 formatted) — 634 or 563 Gbytes, respectively, of highly available storage in a footprint of less than five square feet (.55 square meters).

Storage-system components

The storage system connects to the server by a small computer system interface (SCSI-2) differential bus to a SCSI-2 adapter in the server. To communicate with the storage system, you can use the ArrayGUIde utility that works over the SCSI interface or you can use a terminal connected to the SP's console port.

A storage system includes these components:

- In each server, one or two SCSI-2 adapters
- Storage-system cabinet (deskside) or chassis (rackmount)
- Three or more disk modules
- One or two SPs; for redundancy and greater flexibility, you can acquire a second SP
- Power supplies (VSCs) — two VSCs for a 30-, 20-, or 7-slot storage system or one VSC for a 10-slot storage system; for redundant power with 30-, 20-, or 10-slot systems, you can acquire an additional VSC (30- and 7-slot systems always ship with a redundant VSC)
- Fan module(s) with redundant components

In addition, to implement storage-system write caching, the following components are required:

- For a 30-, 20-, or 10-slot system: two SPs and a battery-backup unit (BBU) — in addition, note that 5 cache vault disks must be present for write-caching to work.
- For a 7-slot system: a PCMCIA memory card with its own battery.

All the components work together as in the following figures.

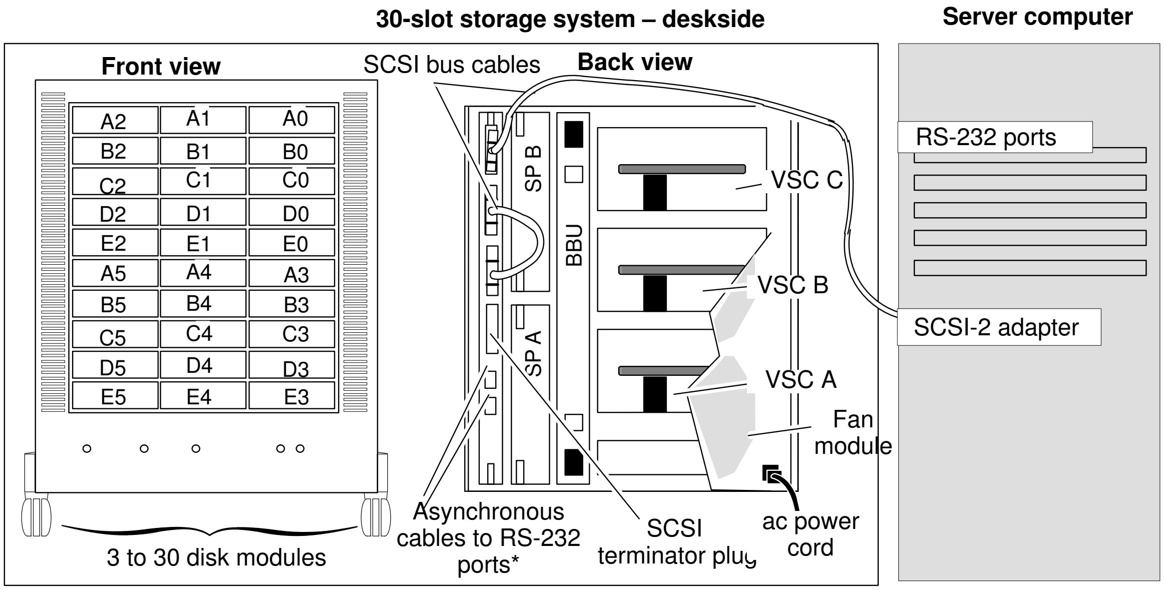

* Asynchronous cables to RS-232 ports are not needed when you use the ArrayGUIde utility.

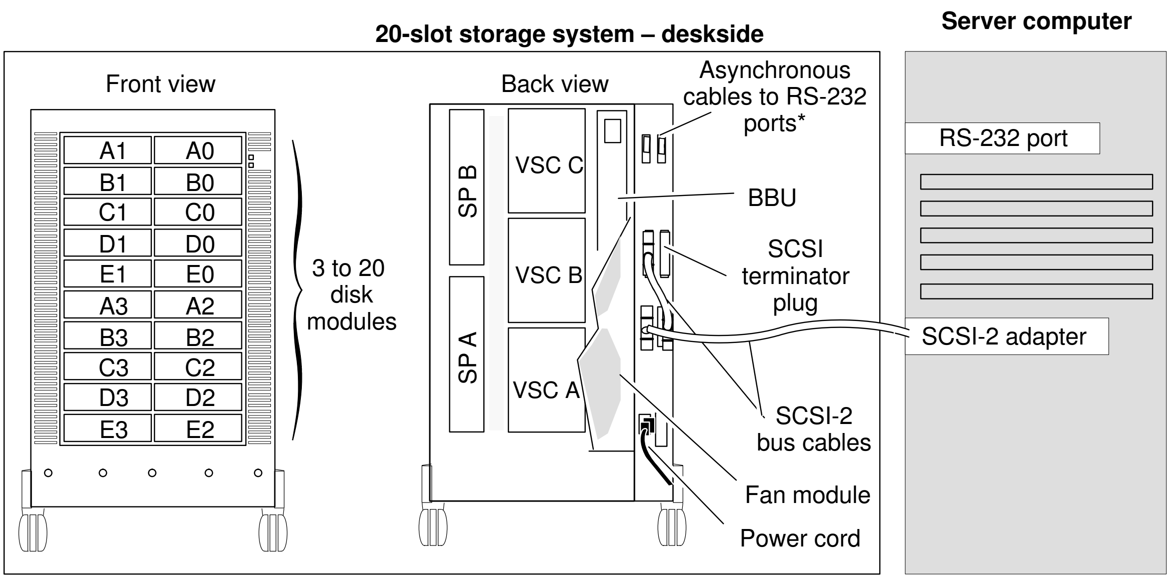

* Asynchronous cables to RS-232 ports are not needed when you use the ArrayGUIde utility.

IMPORTANT Some figures in this guide show the deskside model only. For the rackmount orientation, rotate the deskside figure 90 degrees counterclockwise.

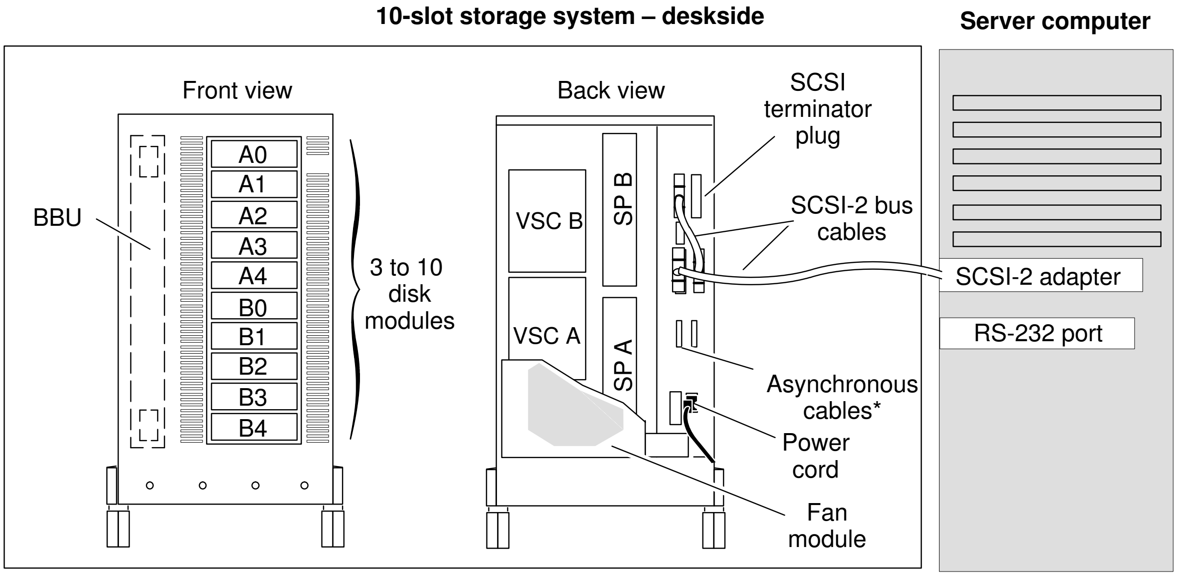

* Asynchronous cables to RS-232 ports are not needed when you use the ArrayGUIde utility.

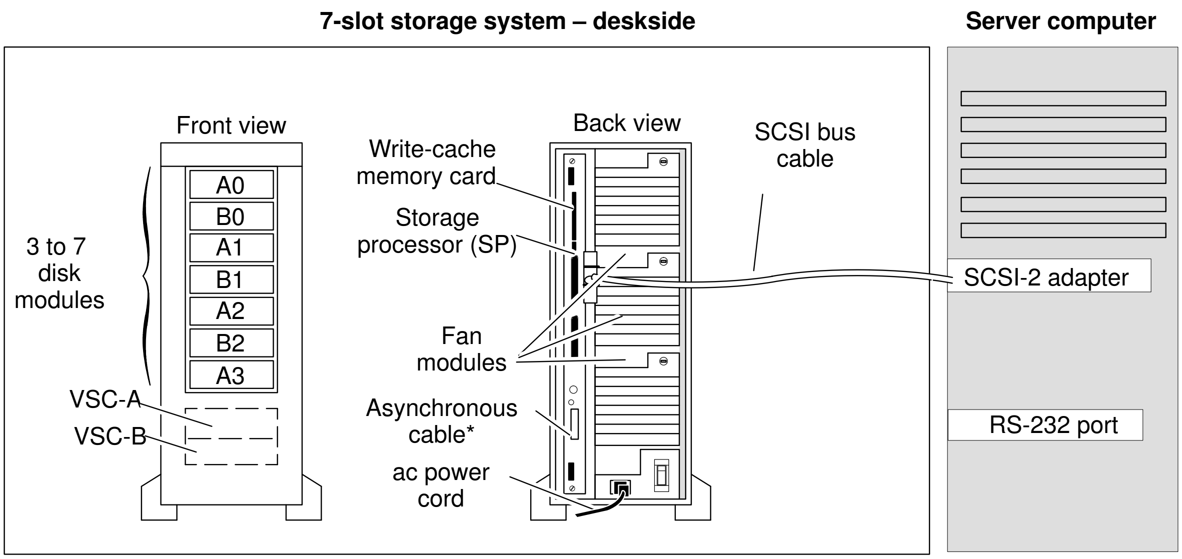

* Asynchronous cables to RS-232 ports are not needed when you use the ArrayGUIde utility.

SCSI-2 adapter

The SCSI-2 adapter (also called a host-bus adapter, HBA) is a printed-circuit board that resides in an I/O slot in the server's cabinet. It transfers data between server memory and the SCSI-2 differential bus. Depending on your computer type, you may have a choice of SCSI-2 adapters.

A SCSI-2 bus cable connects a SCSI-2 adapter to one or two SPs in the storage-system cabinet.

Storage-system chassis (30-slot)

The 30-slot storage-system chassis has slots for disk modules, SPs, VSCs, and the BBU. The disk modules face front; the SP(s), VSCs, BBU, and fan module are accessible from the back. It is available in rackmount only,

The SP directs server I/O into five internal buses, A through E, which have corresponding disk module IDs (A0, B0, C0, and so on). Each internal bus connects to six disk-module slots, as shown next.

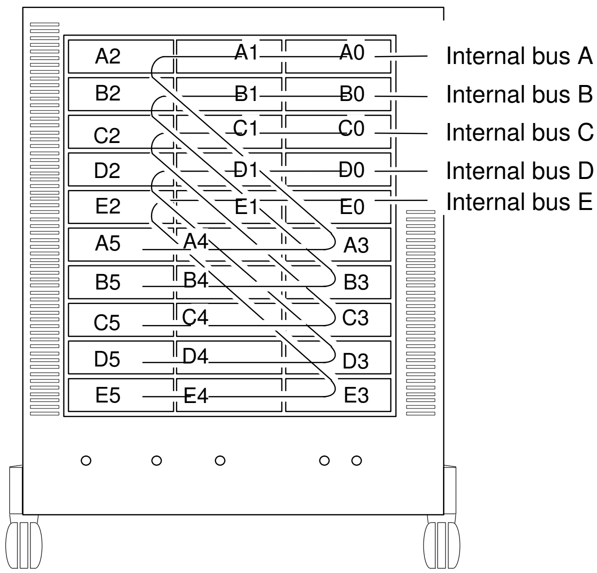

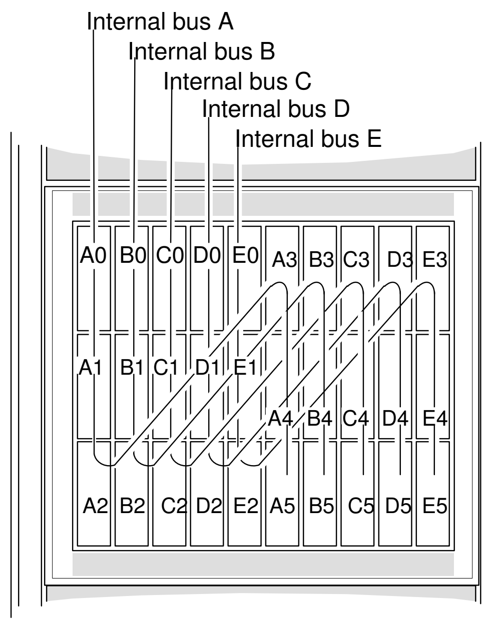

Storage-system chassis (20-slot)

The 20-slot storage-system chassis contains slots for disk modules, SPs, VSCs, and the BBU. The disk modules face front; the SP(s), VSCs, BBU, and the fan module are accessible from the back.

The SP directs server I/O into five internal buses, A through E, which have corresponding disk module IDs (A0, B0, C0, and so on). Each internal bus connects to four disk-module slots, as shown below.

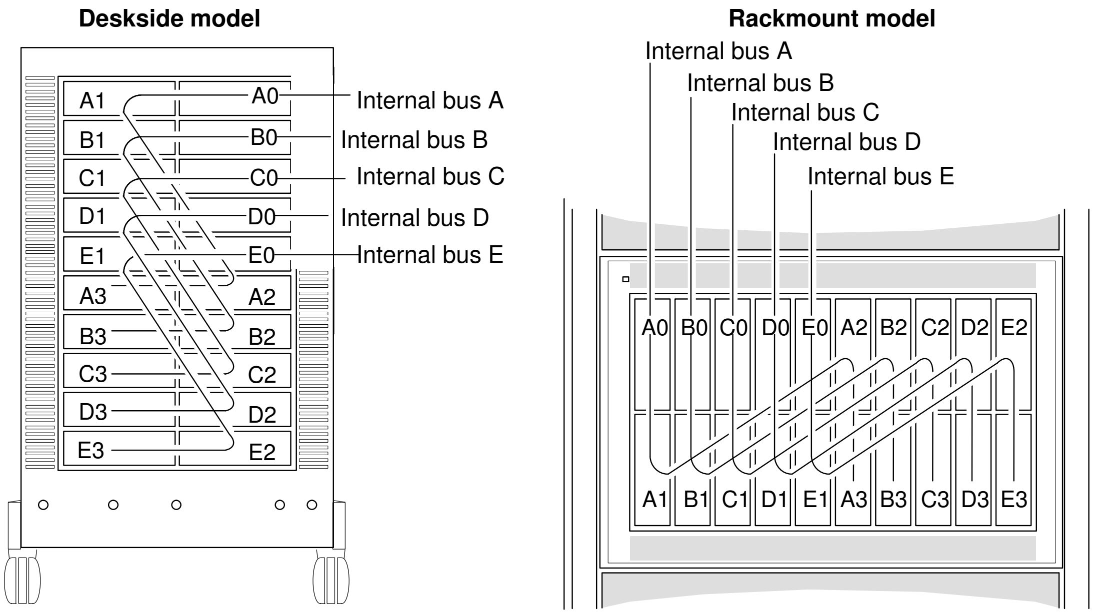

Storage-system chassis (10-slot)

The 10-slot storage-system chassis contains slots for disk modules, SPs, VSCs, and the BBU. The disk modules and BBU face front; the SPs, VSCs, and the fan module are accessible from the back.

The SP directs server I/O into two internal buses, A and B, which have corresponding disk module IDs (A0, B0, and so on). Each internal bus connects to five disk-module slots, as shown below.

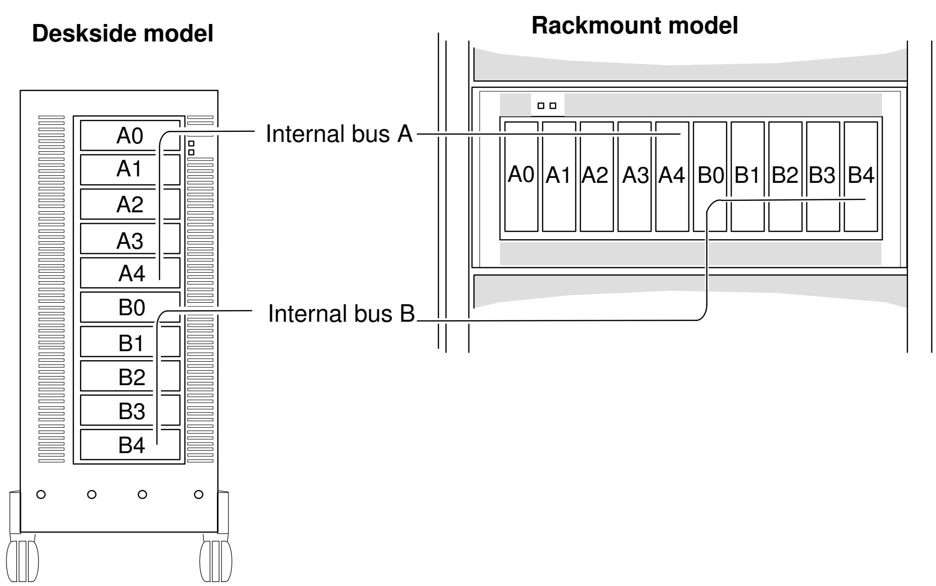

Storage-system chassis (7-slot)

The 7-slot storage-system chassis contains slots for disk modules, SP, and VSCs. The disk modules face front and the VSCs are accessible under the front cover; the SP, PCMCIA memory card, and the three fans are accessible from the back.

The SP directs server I/O into two internal buses, A and B, which have corresponding disk module IDs (A0, B0, and so on). Each internal bus connects to several disk-module slots, as shown below.

Deskside model

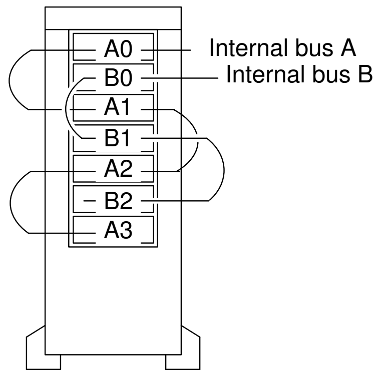

Disk modules

A disk module consists of a disk drive, a power regulator board, internal cabling, and plastic carrier. The carrier has a handle for inserting and removing the module. Three lights on the module indicate status. A label on the carrier shows the module model number and capacity.

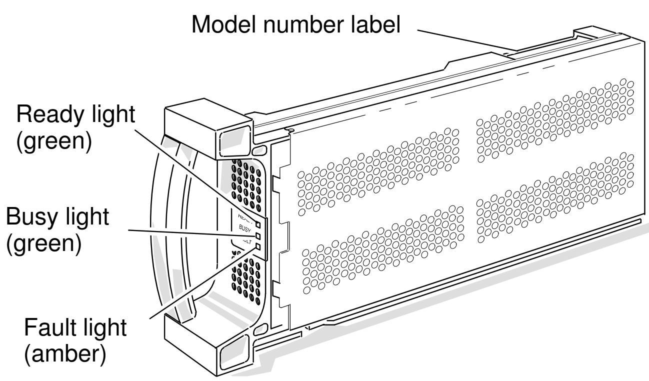

Storage processor (SP)

The SP processes data written to or read from the disk modules, and controls the disk modules in the storage system through a synchronous SCSI-2 bus. An SP in a 30- or 20-slot storage system has five internal buses, each supporting as many as six disk modules for a total of 30 disk modules. An SP in a 10-slot or 7-slot storage system has two internal buses, each supporting as many as five disk modules.

The SP consists of a printed-circuit board with two or four memory modules (SIMMs), status lights, and switches. You can replace a failed SP under power. An SP for a 30- or 20-slot storage system looks like this:

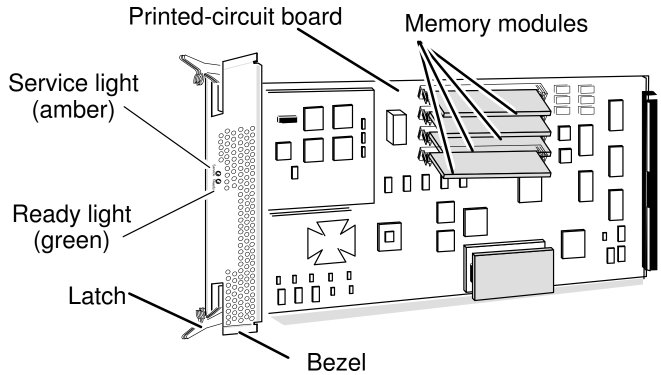

With a specific minimum amount of SP memory, the storage system can support storage-system caching. Storage-system caching enhances performance.

For high availability, all but 7-slot models can support a second SP. A second SP provides a second route to the 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

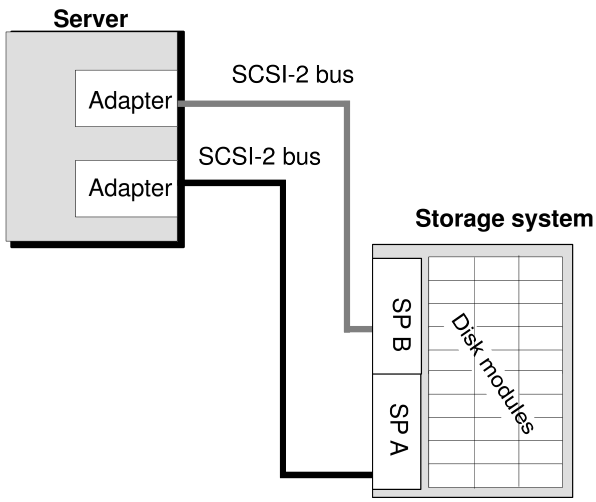

Storage system with two SPs connected to different servers

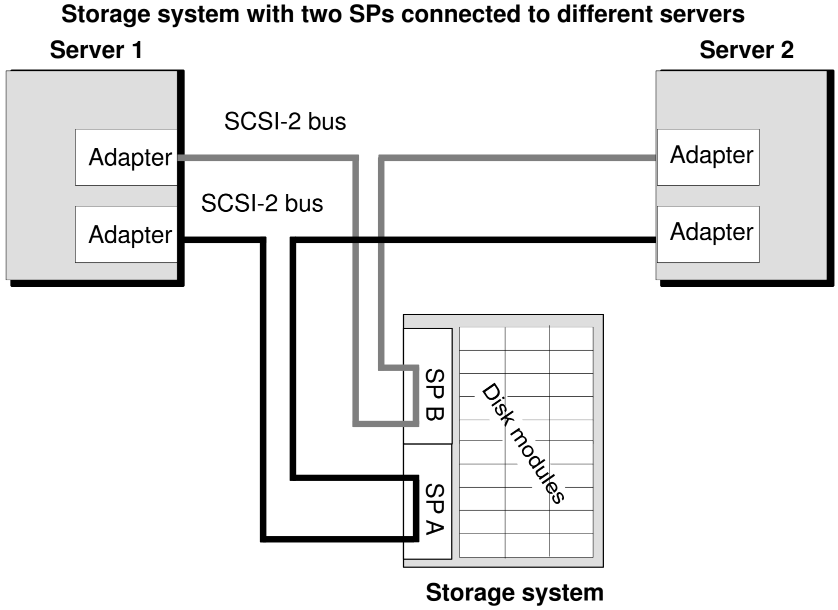

Either SP can control any LUN in the storage system, but only one SP at a time can control a LUN. If one SP cannot access a LUN it controlled (because of a failure), you can transfer control of the LUN to the other SP.

The SP interface programs that let you communicate with the SP are explained on page 16.

Power supplies (VSCs)

The VSCs are power supplies that convert the installation site's ac line voltage to dc voltage required to power the storage-system modules. A 30-slot or 20-slot storage system contains two or three VSCs; three VSCs provide high-availability operation. A 10-slot storage system contains one or two VSCs and a 7-slot storage system contains two VSCs; two VSCs provide high-availability operation.

On a 30-, 20-, or 10-slot system, the VSCs are visible when you swing open the fan module on the back. On a 7-slot system, the VSCs are visible when you remove the front cover (no tools required). You can remove or install a VSC while the storage system is running. A typical VSC looks like this:

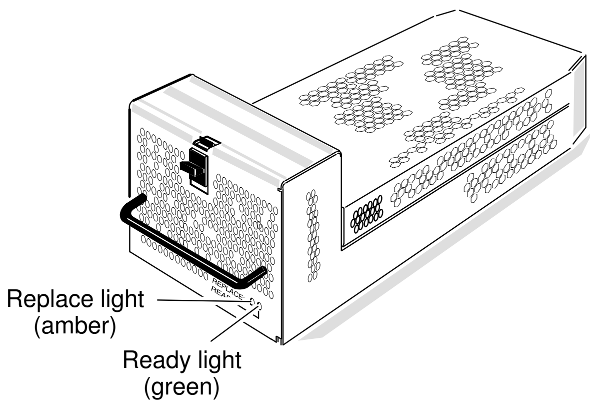

Fan module

All storage systems except 7-slot models use a single fan module, also called a fan pack, that attaches to the back of the chassis. The fan module consists of a control/monitor board and high-capacity, multiple-speed fans — nine fans for a 30-slot system, six fans for a 20-slot system, or three fans for a 7- or 10-slot system. The 7-slot system has separate fan modules, each with its own fan. The following figure shows both the fan module for a 20-slot rackmount chassis (partially opened) and all the fan modules on a 7-slot desktop chassis.

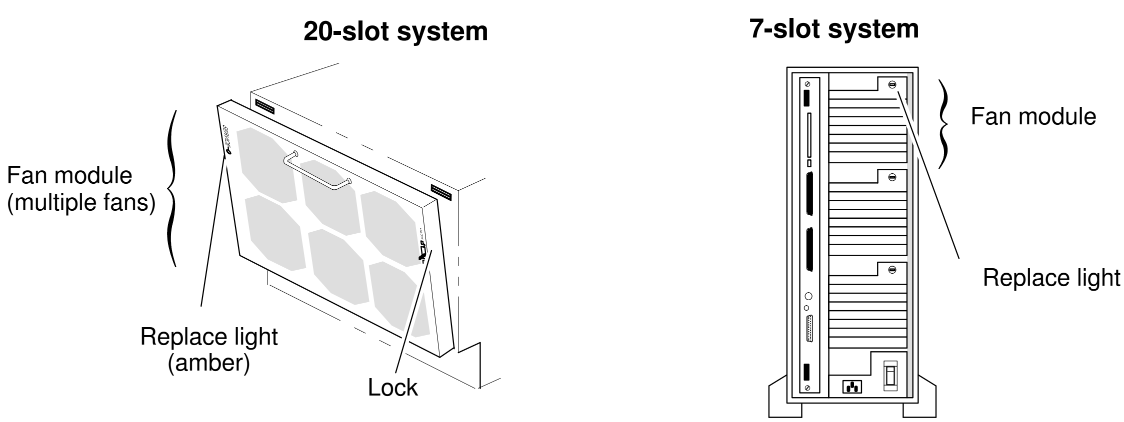

The storage system can continue running after one fan or failure. You can replace a failed fan module while the storage system is running.

Battery backup unit (BBU)

The BBU is required for write caching on 30-, 20-, and 10-slot systems. If line power fails, the BBU lets the SP write cached data to disk and shut down the storage system in an orderly way, as explained in the section on write cache integrity, page 5. When power returns, the BBU recharges its cells automatically. A BBU looks like this:

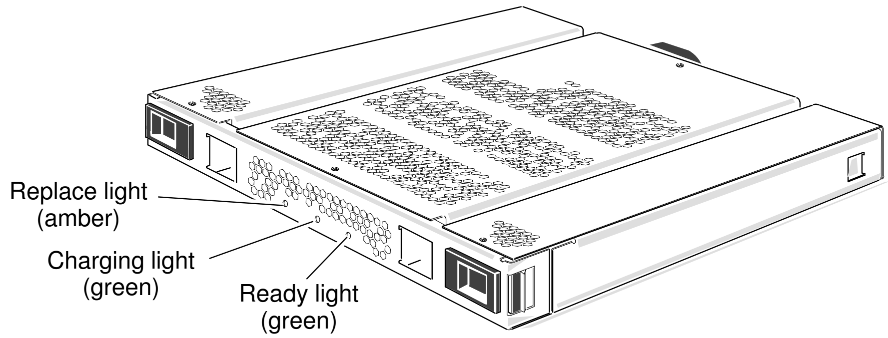

You can install or remove a BBU while the storage system is running.

PCMCIA write-cache memory card

A PCMCIA write-cache memory card is required for write caching on a 7-slot storage system. This card, which is about the size of a credit card, mirrors the SP write cache and has its own replaceable backup battery. The battery maintains the state of PCMCIA card memory during a power or SP failure, as explained in the section on write cache integrity, page 5.

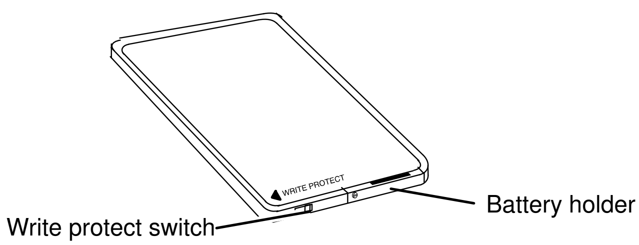

Communication with a storage system

There are several methods to communicate with a storage system:

- ArrayGUIde manager utility, which uses the storage-system SCSI interface, and provides a graphical user interface (GUI) that runs on a server
- direct RS-232 connection to the storage-system console port

Each method lets you bind and unbind disks, examine storage-system status and logged events, and transfer control from one SP to another.

ArrayGUIDE utility

The ArrayGUIde utility has a GUI that runs on Xwindows and lets you manage storage systems connected to a local server or to networked servers. The GUI communicates with an agent program that in turn communicates with the storage system over the SCSI-2 bus.

Direct connection to storage-system console port

The DASSMGR (Disk Array Storage System Manager) interface, which is provided with the storage system, connects to a storage-system console port through the server's RS-232 console port or a telnet connection. Using the RS-232 port requires asynchronous cables.

Server and storage-system configurations

Generally, you will plan the server and storage-system configuration before you plan the disk configuration. There are six storage-system configurations: three for a single server and three for dual servers.

There can be multiple storage systems in any configuration. Usually, with multiple storage systems, you connect the SP As on one bus and the SP Bs on another bus, as shown in the cabling illustration on page NO TAG.

> **IMPORTANT** You must not mix DAS storage systems with other devices, such as CD-ROM or tape drives, on the same SCSI bus. Mixing devices will greatly reduce performance.

A 7-slot storage system has just one SP, so you can use it in configurations that feature a single SP only.

Single-server storage-system configurations

Three single-server storage-system configurations are available for 30-, 20- and 10-slot storage systems. The three are called basic, dual-SP, and dual-adapter/dual-SP. The following table lists the hardware components and features of each single-server configuration. Configuration pictures and a table that shows error recovery features follow this table.

Dual-server configurations are explained in the next section.

Storage-system configurations for a single server

Configuration type	Features
Basic One server One SCSI-2 adapter One SCSI-2 bus One SP	With a RAID group of any level other than 0, applications can continue after failure of any disk module, but cannot continue after failure of SCSI-2 adapter, SCSI-2 bus, or SP.
Dual-SP One server One SCSI-2 adapter One SCSI-2 bus Two SPs	Provides the same features as the basic configuration, but with better performance and higher availability. If one SP fails, I/O requests can be transferred to the other SP.
Dual-adapter/dual-SP One server Two SCSI-2 adapters Two SCSI-2 buses 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 module fails. If one SCSI-2 adapter, SCSI-2 bus, or SP fails, system operator can transfer control to other SCSI-2 adapter.

Basic configuration

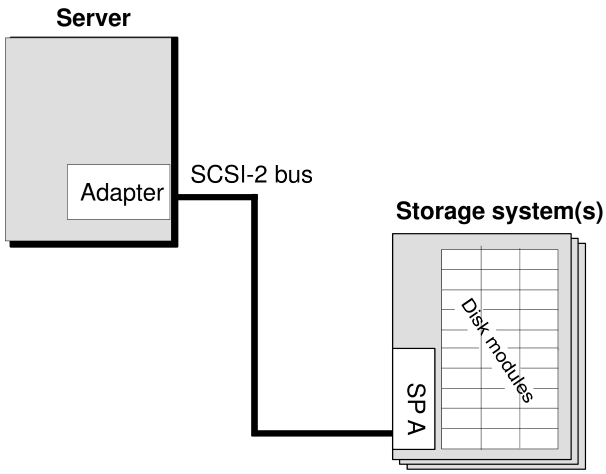

Dual-SP configuration

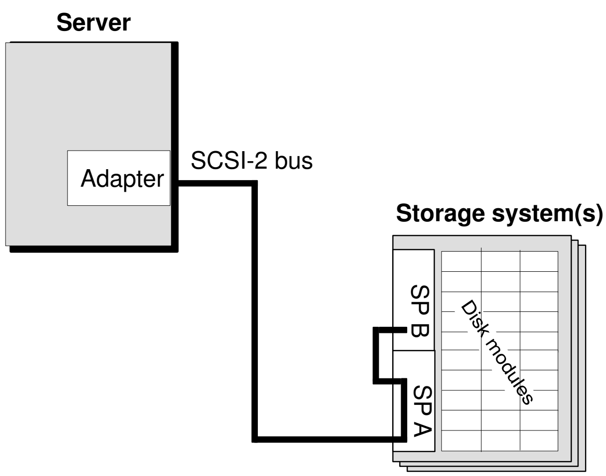

Dual-adapter/dual-SP configuration

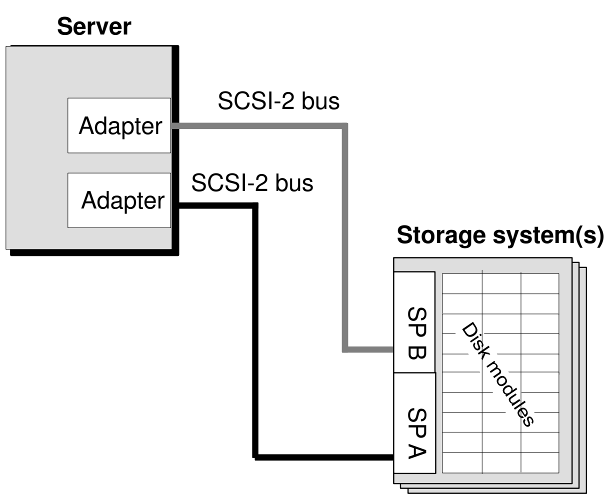

Error recovery features of single-server configurations

Failing component	Configuration	Continue after failure?	What happens and how to recover
Disk module	Basic	Yes	With RAID of any level other than 0, applications continue running; system operator replaces module. The storage system rebuilds the RAID group automatically.
	Dual-SP		
	Dual-adapter/ dual-SP		
SP	Basic	No	Storage system fails; system operator replaces SP and restarts operating system.
	Dual-SP	Yes (with ATF[1])	I/O operations fail to disks owned by failing SP. Application-transparent failover (ATF) automatically transfers control of the failed SP's disk units to the working SP. Applications continue running without receiving any indication of an error. Without ATF, applications fail; operator initiates failover. When convenient, operator can replace the SP and transfer control of disk units to the replacement SP.
	Dual-adapter/ dual-SP		
Fan	Basic	Yes	Applications continue running; system operator replaces fan module at earliest convenience.
	Dual-SP		
	Dual-adapter/ dual-SP		
VSC module (power supply)	Basic	Yes	If the storage system has a redundant VSC module, applications continue running; otherwise, storage system fails. System operator replaces VSC at earliest convenience.
	Dual-SP		
	Dual-adapter/ dual-SP		
SCSI-2 adapter	Basic	No	I/O with storage system fails; authorized service provider replaces adapter; and system operator restarts operating system and applications.
	Dual-SP		
	Dual-adapter/ dual-SP	Yes (with ATF[1])	I/O fails to disks owned by the SP connected to the failing adapter. ATF automatically transfers control of the failed adapter's disk units to the working adapter's SP. Applications continue running without receiving any indication of an error. Without ATF, applications fail; operator initiates failover. When convenient, authorized service provider replaces adapter, and system operator transfers control of disk units to the replacement adapter's SP.
SCSI-2 cable	Basic	No	I/O fails to disks owned by the SP connected to the failing cable. System operator replaces cable and restarts operating system.
	Dual-SP		
	Dual-adapter/ dual-SP	Yes (with ATF[1])	Same as SCSI-2 adapter, except that component is a cable and system operator (not provider) replaces it.

[1] ATF is available as an option.

Dual-server storage-system configurations

With two servers, the storage system offers higher availability: if one server fails, another server can take over and run applications. And sharing a storage system between servers can lower costs. You can connect a second server to an single SP using a Y connector.

Three dual-server storage-system configurations are available. They are called split-bus, dual-initiator, and dual-initiator/dual-bus. The following table lists the hardware components and features of each configuration. Configuration pictures and a summary of error recovery features follow.

Storage-system configurations for dual servers

Configuration type	Features
Split-bus Two servers Two SCSI-2 adapters (one per server) Two SCSI-2 buses (one per server) Two SPs	Resembles two basic configurations using the same storage system, and provide high availability for two servers. Each server and its applications can continue after failure of any disk module. The server using the failed SCSI-2 adapter or SP cannot continue after failure, but the other server can continue and run the other server's applications. Some operating systems support failover software (such as HACMP for AIX) 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 Two servers Two SCSI-2 adapters (one per server) One SCSI-2 bus Two SPs	Two servers share a bus. With a RAID group of any level other than 0, applications can continue after any disk module fails. If a server, SCSI-2 adapter, 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 such as HACMP for AIX.
Dual-initiator/dual-bus Two servers Four SCSI-2 adapters (two per server) Two SCSI-2 buses Two SPs	Two servers share two buses. This provides the highest availability and best storage-system performance for dual-server configurations, since it can handle SCSI cable failure. With a RAID group of any level other than 0, applications can continue after any disk module fails. If a server, SCSI-2 adapter, SCSI-2 bus, 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 such as application–transparent failover (ATF[1]) or HACMP for AIX.

[1] ATF is available as an option.

Split-bus configuration

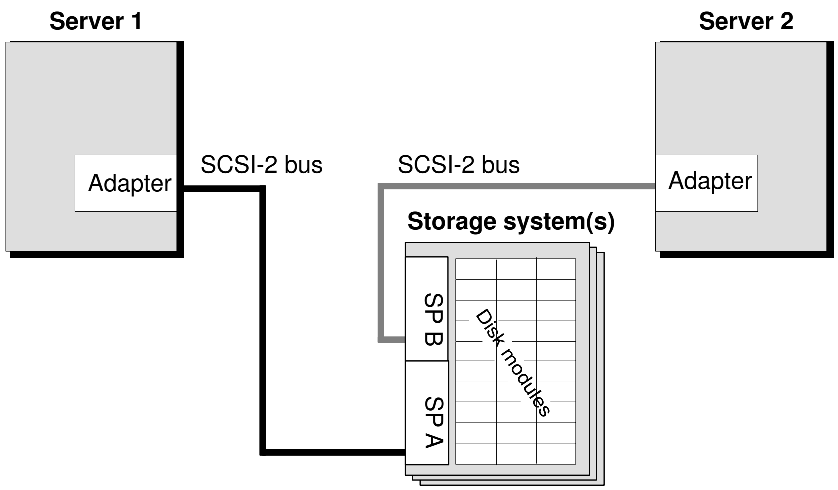

Dual-initiator configuration

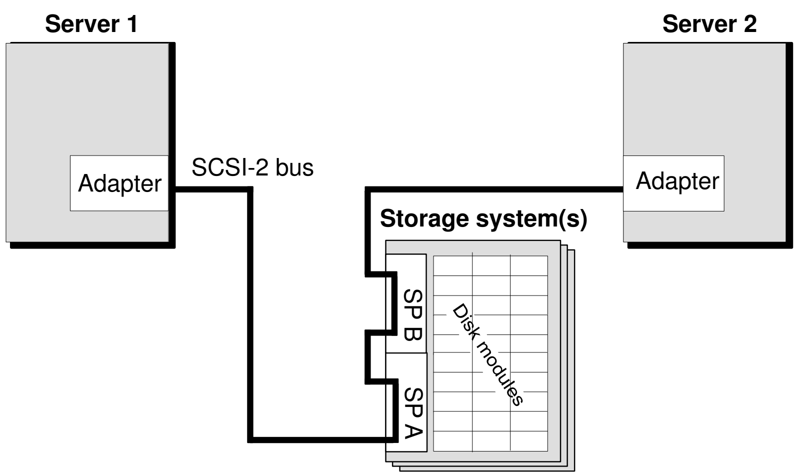

Dual-initiator/dual-bus configuration

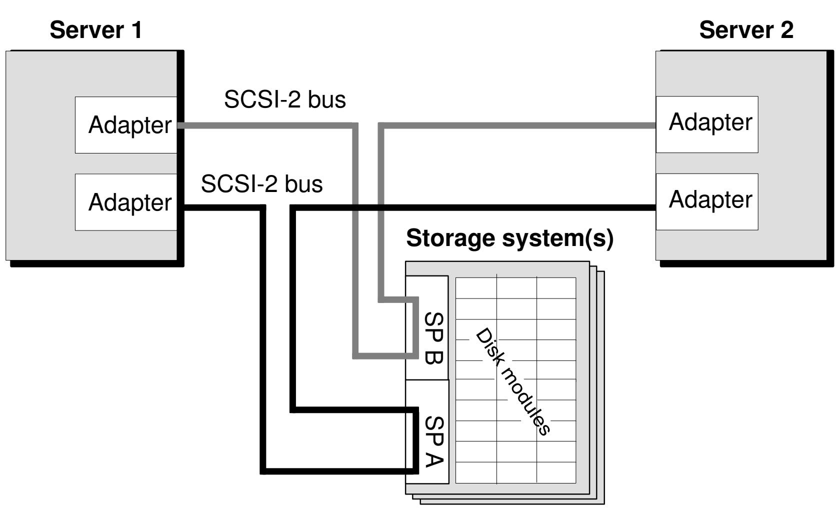

Error recovery features of dual-server configurations

Failing component	Configuration	Continue after failure?	What happens and how to recover
Disk module	Split-bus	Yes	With RAID of any level other than 0, applications continue running; system operator replaces module. The storage system rebuilds the RAID group automatically.
	Dual-initiator		
	Dual-initiator/ dual-bus		
SP	Split-bus	Yes (with failover software[1])	Failover can work as in other dual-server configurations.
	Dual-initiator		I/O fails to disk units owned by failing SP. With failover software, scripts transfer disks and restart applications on other server. Without failover software, applications on failing SP's disks fail. Operator can transfer control of the affected disks to other server or replace the SP, and then restart applications. If control was transferred, later the operator can transfer control of disk units to the replacement SP.
	Dual-initiator/ dual-bus		Same as dual-initiator above, except that applications continue running on other server; they are not interrupted and do not need to be restarted.
Fan	Split-bus	Yes	Applications continue running; system operator replaces fan module at earliest convenience.
	Dual-initiator		
	Dual-adapter/ dual-SP		
VSC module (power supply)	Split-bus	Yes	If the storage system has a redundant VSC module, applications continue running; otherwise, storage system fails. System operator replaces VSC at earliest convenience.
	Dual-initiator		
	Dual-initiator/ dual-bus		
SCSI-2 adapter	Split-bus	Yes (with failover software[1])	Failover can work as in other dual-server configurations.
	Dual-initiator		I/O fails to disk units owned the failed adapter's SP. With failover software, scripts transfer disks and restart applications on other server. Without failover software, applications on failed adapter's SP's disks fail. Operator can transfer control of the affected disks to other server and then restart applications. In either case, when convenient, an authorized service provider replaces the adapter. The system operator can transfer control of disk units to replacement adapter's SP in the original server.
	Dual-initiator/ dual-bus		Same as dual-initiator above, except that applications continue running on other server; they are not interrupted and do not need to be restarted.

[1] Both servers must run the same operating system and failover software that can transfer control if a device fails; for example, ATF or HACMP for AIX. ATF is available as an option.

continues

Error recovery features of dual-server configurations

Failing component	Configuration	Continue after failure?	What happens and how to recover
SCSI-2 cable	Split-bus	Yes (with failover software[1])	Failover can work as in the other dual-server configurations.
	Dual-initiator	No	I/O fails with storage system; operator replaces cable, and restarts operating system and applications.
	Dual-initiator/ dual-bus	Yes (with failover software[1])	I/O fails to disk units owned the failed adapter's SP. With failover software, scripts transfer disks to other server and applications continue. Without failover software, applications on the failing cable's SP's disks fail. The operator can transfer control of the affected disks to the other server or replace the cable, and then restart applications.
Server	Split-bus	Yes (with failover software[1])	Failover can work as in the other dual-server configurations.
	Dual-initiator		With failover software, scripts transfer disks and restart applications on the other server. Without failover software, applications on the failing server's disks fail and remain down until the server is fixed. The operator can transfer control of the affected disks to the other server, and then restart applications. Authorized service provider fixes server.
	Dual-initiator/ dual-bus		

1 Both servers must run the same operating system and failover software that can transfer control if a device fails; for example, ATF or HACMP for AIX. ATF is available as an option.

concluded

Disk configuration definitions

This section defines disk configuration terms and the types of disk configuration you can create in a disk–array storage system.

The disk–array storage system uses RAID (redundant array of independent disks) technology. RAID technology groups separate disk modules into one physical storage unit 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 disk modules into one physical disk unit by *binding* them using a storage–system management utility.

Four of the RAID groups defined in this section use *disk striping* and two use *mirroring*. These two terms have the following meanings.

Disk striping

Using disk stripes, the storage–system hardware can read from and write to multiple disk modules 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 or written to each module 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. More details on stripe size appear in the description of RAID groups later.

The storage system uses disk striping with most types of RAID group.

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 AIX operating system synchronizes the images.

With the storage system, you can create a hardware mirror by binding disk modules as a RAID–1 mirrored pair or a RAID–1/0 group (a mirrored RAID–0 group); the hardware will then mirror the modules automatically. Or you can use AIX software mirroring with RAID–0 groups or individual units that have no inherent data redundancy. With software mirroring, the AIX operating system mirrors the images. RAID–1 mirrored pairs and RAID–1/0 groups are further explained in the following sections.

Disk configuration types

Types of disk configuration you can create are RAID–5 group, RAID–3 group, RAID–1 mirrored pair, RAID–0 group, RAID–1/0 group, individual disk unit, and hot spare.

RAID–5 group (individual access array)

A RAID–5 group usually consists of five modules (but can have three to sixteen). A RAID–5 group uses disk striping.

The storage system writes parity information that lets the group continue operating if a disk module fails. When you replace the failed module, 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 module.

For highest availability for a RAID–5 group in a 30– or 20–slot storage system, place the disk modules on different SCSI internal buses (A, B, C, D, or E). On a 10–slot or 7–slot storage system, there are only two internal buses, so it does not matter which bus you bind them on.

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

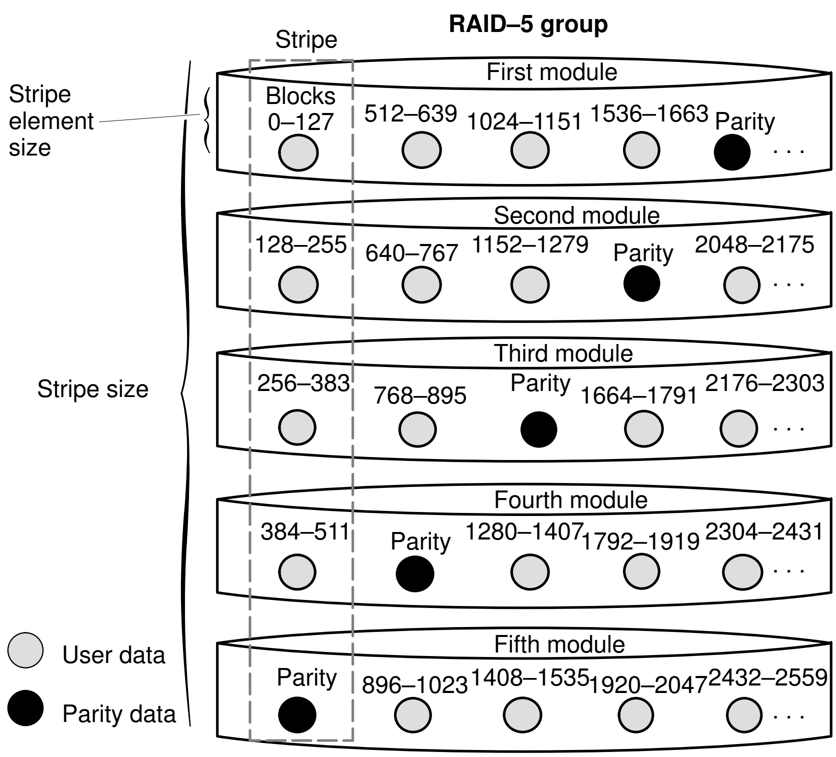

The storage system performs more steps when writing data to a RAID–5 group than to any other type of group. For each write, the storage system

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 disk modules. The hardware always reads from or writes to all the disk modules. A RAID–3 group uses disk striping. A RAID–3 group requires five disk modules on five different internal SCSI buses (A, B, C, D, and E).

The storage system writes parity information that lets the group continue operating if a disk module fails. When you replace the failed module, 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 module.

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 module to the second, third, and fourth, then the first, and so on.

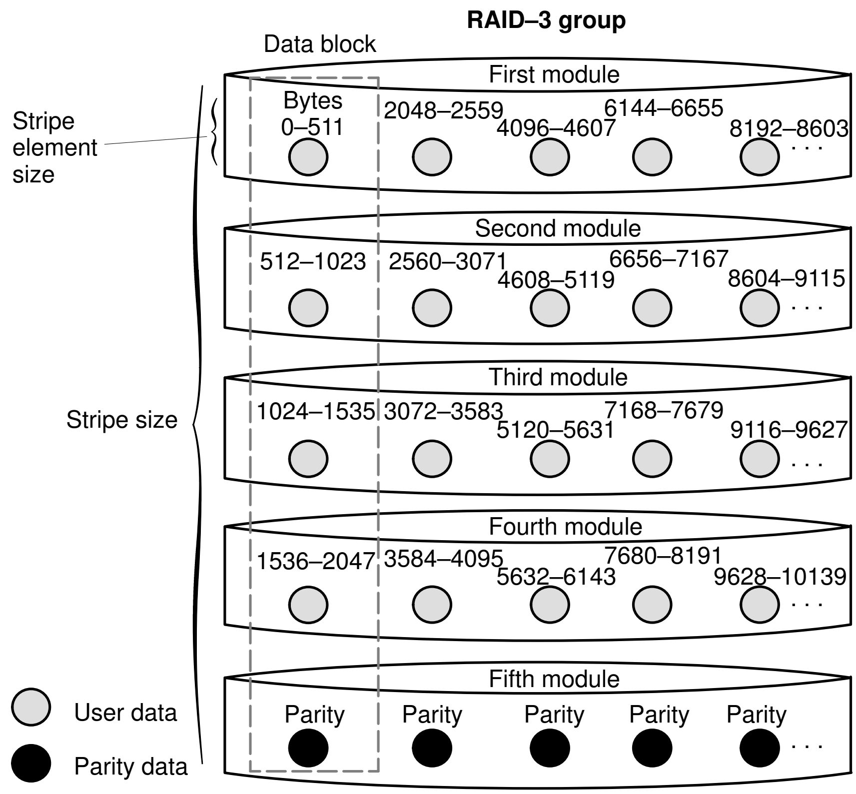

RAID–3 differs from RAID–5 in several important ways. First, in a RAID–3 group the hardware processes disk requests serially, whereas in a RAID–5 group the hardware can interleave disk requests. Second, with a RAID–3 group, the parity information is stored on one disk module; with a RAID–5 group, it is stored on all modules. Finally, with a RAID–3 group, the I/O occurs in small units (one sector) to each disk module. A RAID–3 group works well for single–task applications that use I/Os of blocks larger than 64 Kbytes. To reach the RAID–3 performance potential, the blocks must be aligned to start at disk addresses that are multiples of 2 Kbytes from the beginning of the logical disk.

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 other types of RAID groups, since SP resources (including memory) are best devoted to the RAID–3 groups. RAID–3 is available with 30– and 20–slot systems only.

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

1. Calculates the parity data
2. Writes the new user and parity data. (If the write is not a multiple of 2 Kbytes or the starting disk address of the I/O does not begin at an even 2–Kbyte boundary from the beginning of the logical disk, all the steps required for a RAID–5 write are needed.)

RAID–1 mirrored pair

A RAID–1 group consists of two disk modules which 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 disk modules in such a mirror separately, you must unbind the mirror (losing all data on it), rebind the module in the configuration you want, and software format the physical disk unit.

With a disk–array storage system, RAID–1 hardware mirroring has the following advantages: automatic operation (you do not have to issue any commands to initiate it); physical duplication of images; and 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 disk modules, as follows.

RAID–1 mirrored pair

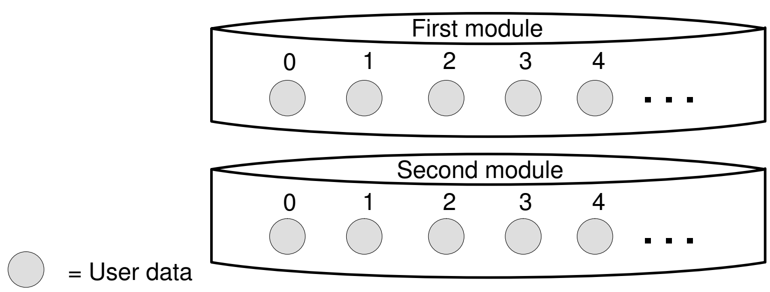

For the most highly available mirror of any type, we recommend that you configure the mirror with each disk module on a different internal SCSI bus.

RAID–0 group (nonredundant array)

A RAID–0 group consists of three to a maximum of sixteen disk modules. A RAID–0 group uses disk striping, in which the hardware writes to or reads from multiple disk modules simultaneously. For more information, see the section on disk striping, page 24.

Unlike the other RAID levels, with RAID–0 the hardware does not maintain parity information on any disk module; this type of group has no inherent data redundancy. RAID–0 offers enhanced performance through simultaneous I/O to different disk modules.

You can use AIX software mirroring with the RAID–0 group to provide high availability. A desirable alternative to RAID–0 is RAID–1/0, next.

RAID–1/0 group (mirrored RAID–0 group)

A RAID–1/0 group consists of four, six, eight, ten, twelve, fourteen, or sixteen disk modules. These disk modules make up two mirror images, with each image including two to five disk modules. In this configuration, the hardware automatically mirrors the modules. A RAID–1/0 group uses disk striping. It combines the speed advantage of RAID–0 with the redundancy advantage of mirroring.

The following figure shows the distribution of user data with the default stripe element size of 128 sectors (65,536 bytes) in a six–module RAID–1/0 group. Notice that the disk block addresses in the stripe proceed sequentially from the first mirrored modules (first and fourth modules) to the second mirrored modules (second and fifth modules), to the third mirrored modules (third and sixth modules), then from the first mirrored modules, and so on.

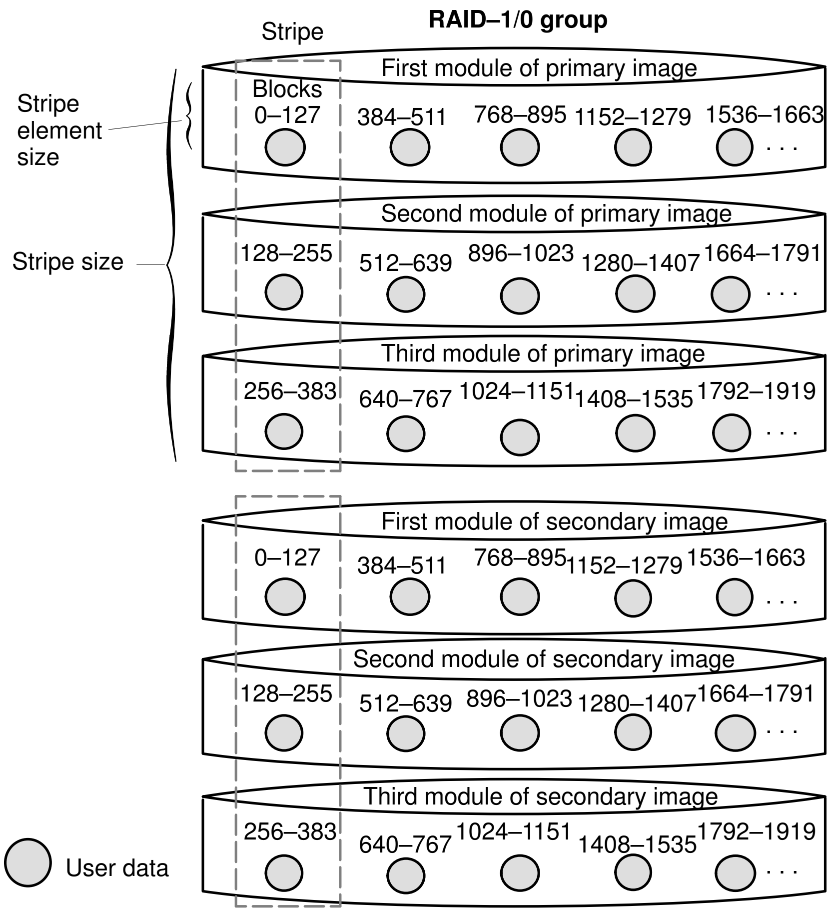

A RAID–1/0 group can survive the failure of multiple disk modules, providing that one module in each image pair survives. So you should bind the modules in each image pair on different internal SCSI buses.

For practical use of disk modules, we recommend RAID–1/0 groups for 30–, 20–, and 10–slot storage systems only.

For a 30– or 20–slot storage system – For highest availability and best performance, place the modules on a different internal SCSI bus, A, B, C, D, and E. For example,

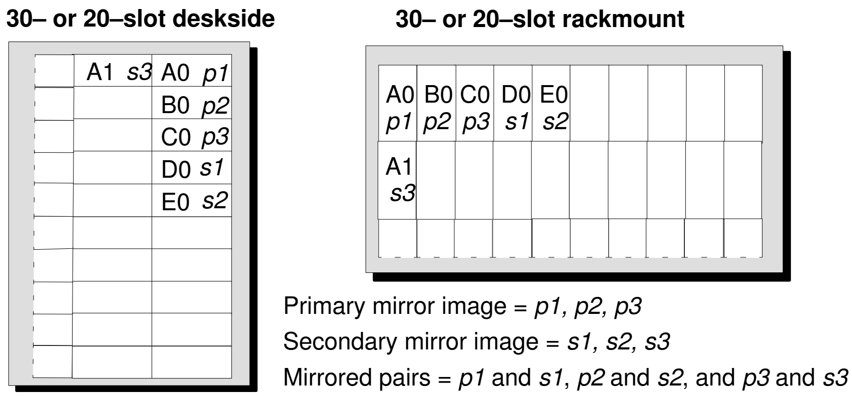

This RAID–1/0 group has six modules that make up the primary and secondary image pairs. Modules *p1* and *s1* are in slots A0 and D0, *p2* and *s2* are in slots B0 and E0, and *p3* and *s3* are in slots C0 and A1. The group can continue after failure of A0 or D0, B0 or E0, and C0 or A1.

For a 10–slot storage system – When you plan a RAID–1/0 group, for highest availability and best performance, place the disk modules for the primary image on one internal SCSI bus and the disk modules for the secondary image on the other internal SCSI bus. For example,

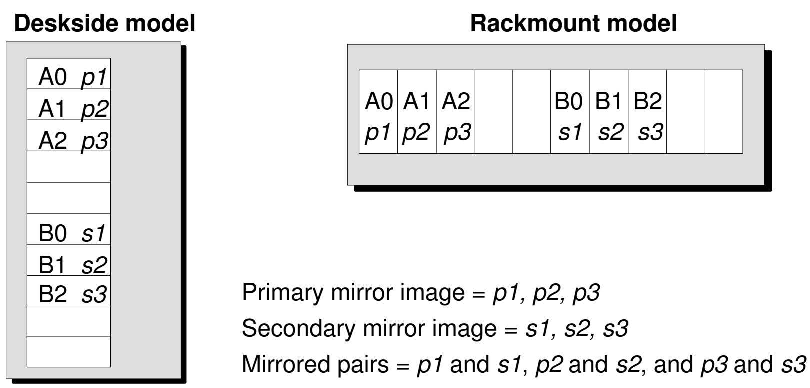

This RAID–1/0 group has six modules that make up the primary and secondary image pairs. Modules *p1* and *s1* are in slots A0 and B0, *p2* and *s2* are in slots A1 and B1, and *p3* and *s3* are in slots A2 and B2. The group can continue after failure of A0 or B0, A1 or B1, and A2 or B2.

Individual disk unit

An individual disk unit is a disk module bound to be independent of any other disk module in the cabinet. An individual unit has no inherent high availability feature, but you can make it highly available by using software mirroring with another individual unit, preferably on a different internal SCSI bus.

Hot spare

A hot spare is a dedicated replacement disk module 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 module'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 starts copying the data from the former hot spare onto the replacement disk.

The SP finishes rebuilding the disk module before it begins copying data, even if you replace the failed disk during the rebuild process.

When the copy is done, the disk group consists of disk modules 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 module has failed, find a suitable replacement module, and insert the module.

> **IMPORTANT** When you plan to use a hot spare, make sure the module has the capacity to serve in any RAID–5, RAID–3, RAID–1, or RAID 1/0 group in the storage system. The storage system cannot use a hot spare that is smaller than the module it replaces.

You can have one or more hot spares per storage system. You can make any module in the storage system a hot spare, except for the following modules:

30–slot:	A0, B0, C0, D0, E0, A3, and A4
20–slot:	A0, B0, C0, D0, E0 and A3
10–slot:	A0, A1, A2, A3, A4, and B0
7–slot:	A0, B0, A1

Examples of hot spare usage for 30–slot, 20–slot, and 10–slot systems follow.

30–slot storage system

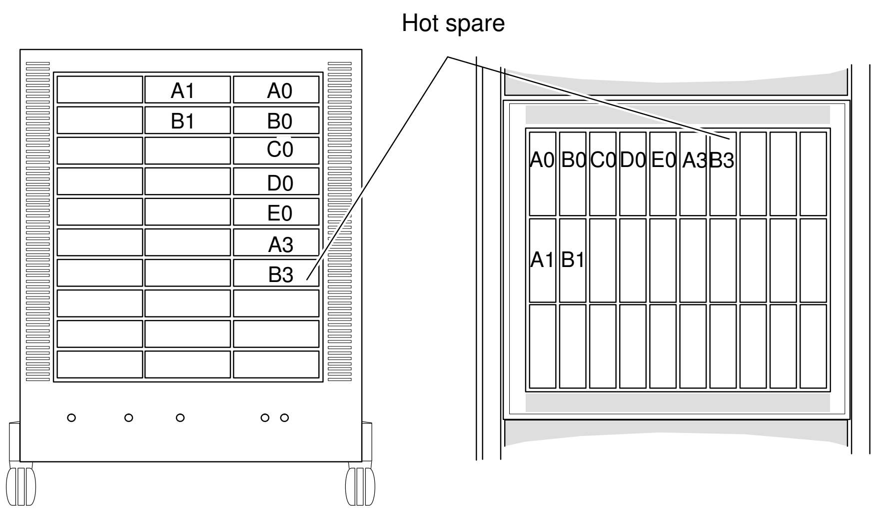

1. RAID–5 group is disk modules A0–E0; RAID–1 mirrored pair is modules A1 and B1; hot spare is module B3.
2. Disk module D0 fails.
3. RAID–5 group becomes A0, B0, C0, B3, and E0; now no hot spare is available.
4. System operator replaces failed module in D0 with a functional module.
5. RAID–5 group once again is A0–E0 and hot spare is in B3.

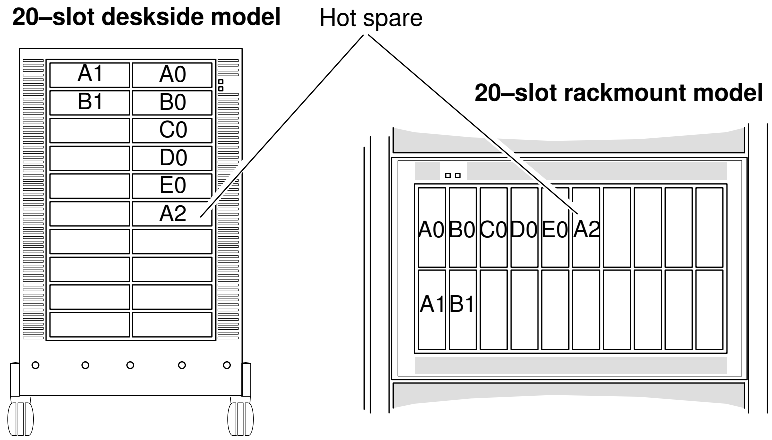

1. RAID–5 group is disk modules A0–E0; RAID–1 mirrored pair is modules A1 and B1; hot spare is A2.
2. Disk module D0 fails.
3. RAID–5 group becomes A0, B0, C0, A2, E0; now no hot spare is available.
4. System operator replaces failed module D0 with a functional module.
5. RAID–5 group once again is A0–E0 and hot spare is A2.

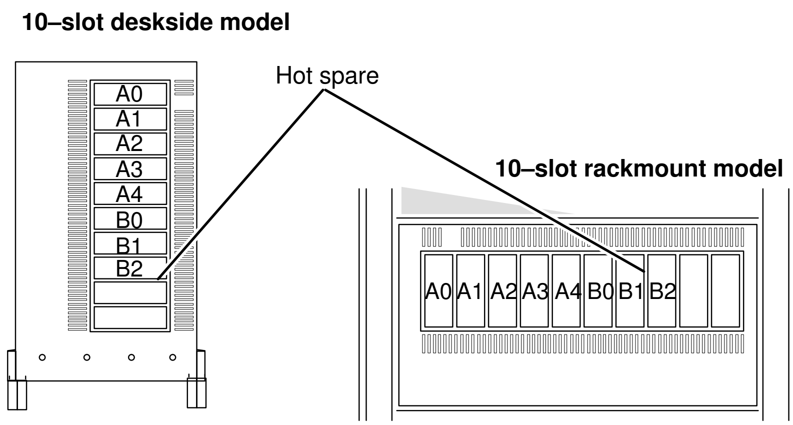

1. RAID–5 group is disk modules A1–A4 and B1; RAID–1 mirrored pair is modules A0 and B0; hot spare is B2.
2. Disk module A3 fails.
3. RAID–5 group is A1, A2, B2, A4, and B1; now no hot spare is available.
4. System operator replaces failed disk module A3 with a functional module.
5. RAID–5 group is once again A1–A4 and B1, and hot spare is B2.

Disk configuration benefits and tradeoffs

This section reviews the disk configuration types and explains their benefits and tradeoffs.

You can create seven types of physical disk unit (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
- RAID–0 group (nonredundant individual access array); no inherent high availability features, but can be software mirrored through the mirroring capability of the AIX operating system
- Individual unit; no inherent high availability feature,s but can be software mirrored through the mirroring capability of the AIX operating system
- Hot spare; serves only as an automatic replacement for any disk module in a RAID configuration other than 0; does not store data during normal system operations

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

Performance

RAID–5, with individual access, provides high read throughput for small block–size requests (blocks of 2 to 8 Kbytes) by allowing simultaneous reads from each disk module 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 storage–system caching, which means you do not need the second SP and BBU that caching requires.

Generally, the performance of a RAID–3 unit increases as the size of the I/O request increases. Read performance increases rapidly with read requests up to 1 Mbyte. Write performance increases greatly for sequential write requests that are greater than 256 Kbytes. With a 30– or 20–slot storage system, 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 with RAID–5 or RAID–1/0.

A **RAID–1** mirrored pair has its disk modules locked in synchronization, but the SP can read data from the module whose read/write heads are closer to it. Therefore RAID–1 read performance can be twice that of an individual disk and write performance 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 disk modules in the group. Since RAID–1/0 locks pairs of RAID–0 disk modules 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 physical disk, use a RAID–1/0 or RAID–0 group. A RAID–1/0 group requires at least six disk modules; a RAID–0 group, at least three modules.

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

Data availability and disk space usage

If data availability is critical and you cannot afford to wait hours to replace a disk module, 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, 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– groups, since only 50% of the total disk capacity is available for user data, as shown on page 34.

A RAID–5 or RAID–3 group provides high data availability, but requires more modules than a mirrored pair. In a RAID–5 or RAID–3 group of five disk modules, 80% 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 physical disk configurations

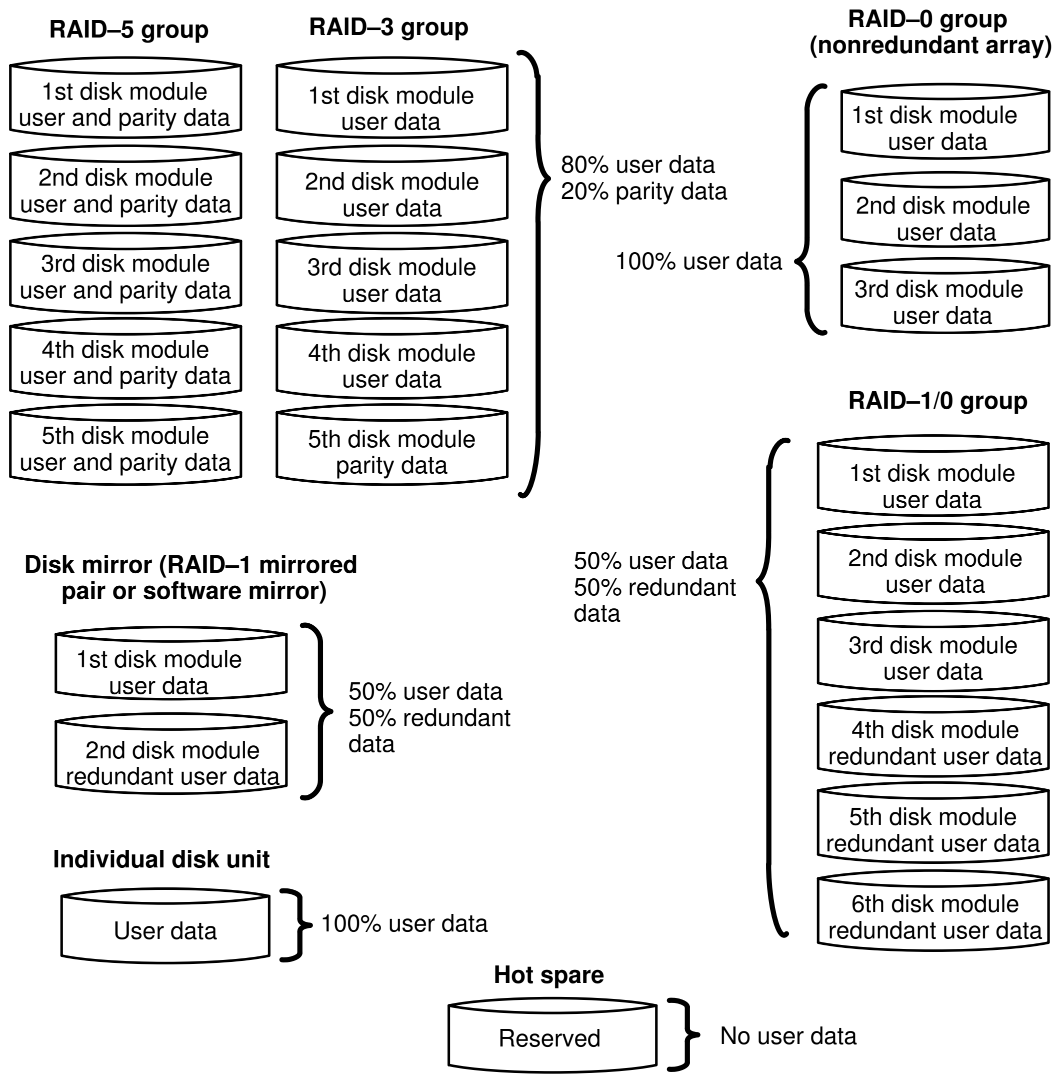

You can use AIX software mirroring for an individual unit or RAID–0 group. For a comparison of software and hardware mirroring, see the section on mirroring, page 24.

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 above. For high availability with an individual disk, you can use software mirroring.

Guidelines for physical disk configurations

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 disk configurations.

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 (RAID–5 write performance can be improved with storage–system caching)

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 beginning of the logical disk.

Use a RAID–1 mirrored pair (or software mirror) 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 module in the group fails
- Minimizing the degraded performance caused by disk module failure in a RAID–5 or RAID–3 group is important

Sample applications for physical disk configurations

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 less–than–normal percentage of write operations (writes are 33% 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 numerous 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 modules 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 mirroring the RAID–0 group using AIX software mirroring. A RAID–0 group (like a RAID–5 group) requires a minimum of three disk modules; software mirroring a RAID–0 group requires the same number of modules 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 disk modules.

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, you can always use the software mirroring function provided by the AIX operating system.

The performance of an individual unit is slightly less than a standard SCSI disk not in a 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 it where you must regain high availability quickly without human intervention if any disk module in such a RAID group fails. A hot spare also minimizes the period of degraded performance after a RAID–5 or RAID–3 module fails.

Dual SPs and routes to disk units

If a storage system has two SPs, you can choose which disks to bind on each SP. The SP on which you bind a disk module is the default owner of that physical disk unit. The route through the SP that owns a unit determines the primary route to that unit and determines the device name of the unit. The route through the other SP is the secondary route to the unit.

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

Sample physical disk configurations

This section shows two storage–system configurations — a single–server, dual–adapter/dual–SP configuration and a dual–server, dual–initiator/dual–bus configuration — with sample physical disk configurations.

Single–server example, 30– or 20–slot

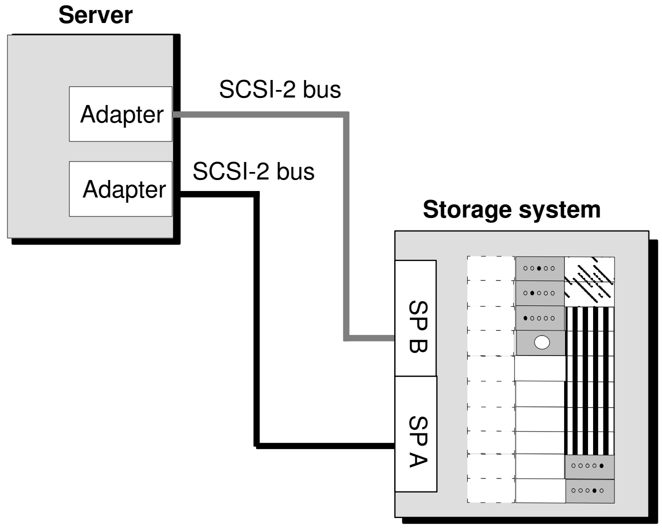

If each disk module shown holds 9 Gbytes, then the storage system provides 81 Gbytes of disk storage with 72 Gbytes highly available. The physical disk configurations are

Unit **S** on two disk modules bound as a RAID-1 mirrored pair for 9 Gbytes of storage; for the user directories.

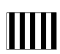

Unit **T** on six modules bound as a RAID-1/0 group (mirrored RAID-0 group) for 27 Gbytes of storage – for clients and mail.

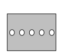

Unit **U** on five modules bound as a RAID-5 group (individual access array) for 36 Gbytes of storage; for the database.

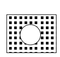

Unit **V** on one module bound as an individual unit for 9 Gbytes of storage; temporary directory for users.

Single–server example, 7–slot

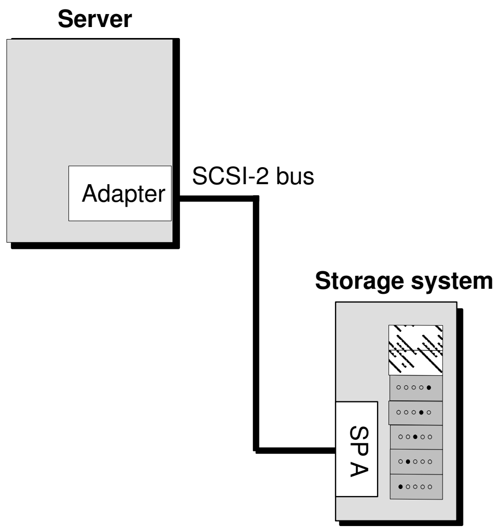

If each disk module shown holds 9 Gbytes, then the storage system provides 45 Gbytes of disk storage, all highly available. The physical disk configurations are

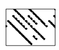

Unit **S** on two disk modules bound as a RAID–1 mirrored pair for 9 Gbytes of storage; for the user directories.

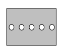

Unit **T** on five disk modules bound as a RAID–5 group (individual access array) for 36 Gbytes of storage; for the database.

Dual–server example, 30– or 20–slot

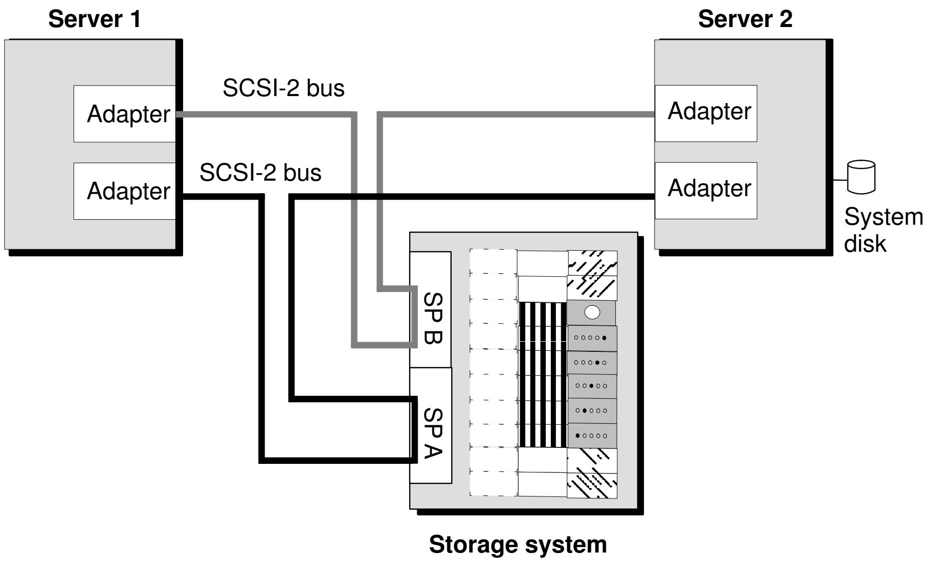

If each disk module shown holds 9 Gbytes, then the storage system provides Server 1 with 54 Gbytes of disk storage, 45 Gbytes highly available; it provides Server 2 with 36 Gbytes of storage, all highly available. Each server has its own SP, which controls that server's storage–system disks; those disks remain primary to that server. The disk configurations are

Server 1 (primary route is SP A; secondary route is SP B)

Unit **S** on two disk modules bound as a RAID–1 mirrored pair for 9 Gbytes of storage; for the user directories.

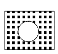

Unit **T** on one disk module bound as an individual unit for 9 Gbytes of storage; for a temporary user directory.

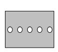

Unit **U** on five disk modules bound as a RAID–5 group (individual access array) for 36 Gbytes; for Server 1's database.

Server 2 (primary route is SP B; secondary route is SP A)

Unit **V** on two disk modules, bound as a RAID–1 mirrored pair, for 9 Gbytes of storage; for user directories.

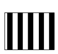

Unit **W** on six disk modules bound as a RAID–1/0 group (mirrored RAID–0 group) for 27 Gbytes of storage; for customer account information.

Site requirements

This section describes storage-system size, power and clearance needs.

30-slot storage system requirements

Rackmount model

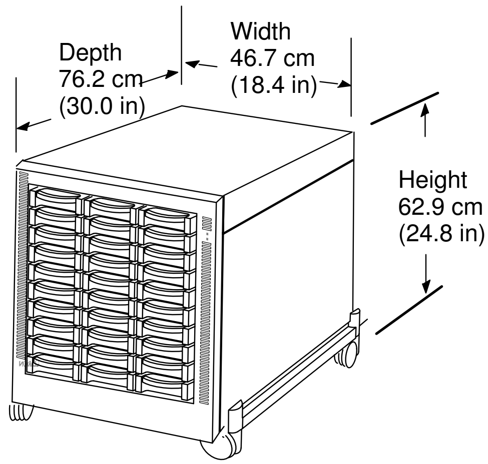

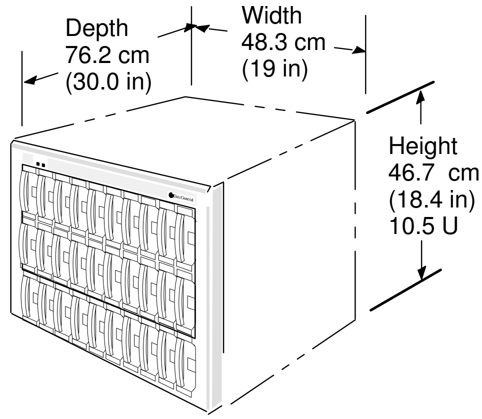

Service clearances	Front: 81.3 cm (32 in) Rear: same

Weight (without packaging)	**Deskside**	**Rackmount**
Maximum (30 disk modules, 2 SPs, 3 VSCs, BBU)	06.6 kg (235 lb)	99 kg (220 lb)

Operating environment	
Temperature	10ºC to 38ºC (50º F to 100º F)
Relative humidity (noncondensing)	20% to 80%
Altitude	2438 m (8000 ft)
Heat dissipation (max)	3570 BTU/hour

Power requirements	
Voltage rating	200 to 240 V ac, single-phase, 47 to 63 Hz; power supplies are auto-sensing and auto-ranging
Power consumption	1000 VA max (apparent), 1050 VA max (true)
Phase	1 (USA), or 2 or 3 (outside USA)

Power cables	
USA primary power	1.8 m (6.0 ft); NEMA L5-15P connection, L5-15 mating connection
Other locations	Local standard ac connection

20-slot storage system requirements

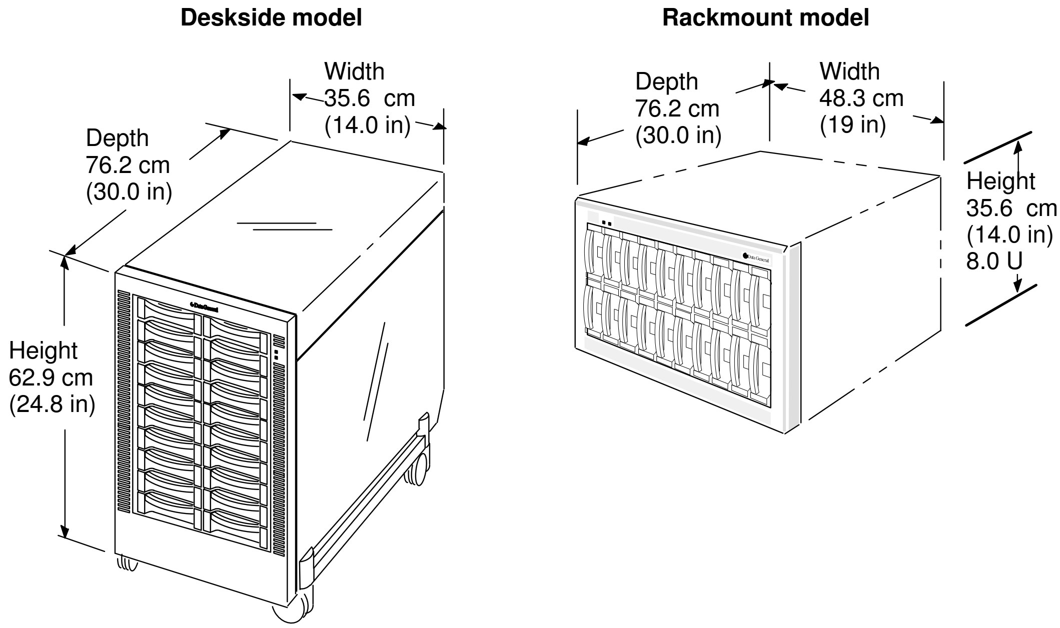

Service clearances	Front: 81.3 cm (32 in) Rear: same	Front: 81.3 cm (32 in) Rear: same

Weight (without packaging)	**Deskside**	**Rackmount**
Maximum (20 disk modules, 2 SPs, 3 VSCs, BBU)	83.9 kg (186.4 lb)	78 kg (173.2 lb)

Operating environment

Temperature	10ºC to 38ºC (50º F to 100º F)
Relative humidity (noncondensing)	20% to 80%
Altitude	2438 m (8000 ft)
Heat dissipation (max)	3168×10^3 Joules, 3000 BTU/hour

Power requirements

Voltage rating	90 V ac to 264 V ac, single-phase, 47 to 63 Hz; power supplies are auto-sensing and auto-ranging
Power consumption	900 VA max (apparent), 880 VA max (true)
Phase	1 (USA), or 2 or 3 (outside USA)

Power cables

USA primary power	1.8 m (6.0 ft); NEMA L5-15P connection, L5-15 mating connection
Other locations	Local standard ac connection

10-slot storage system requirements

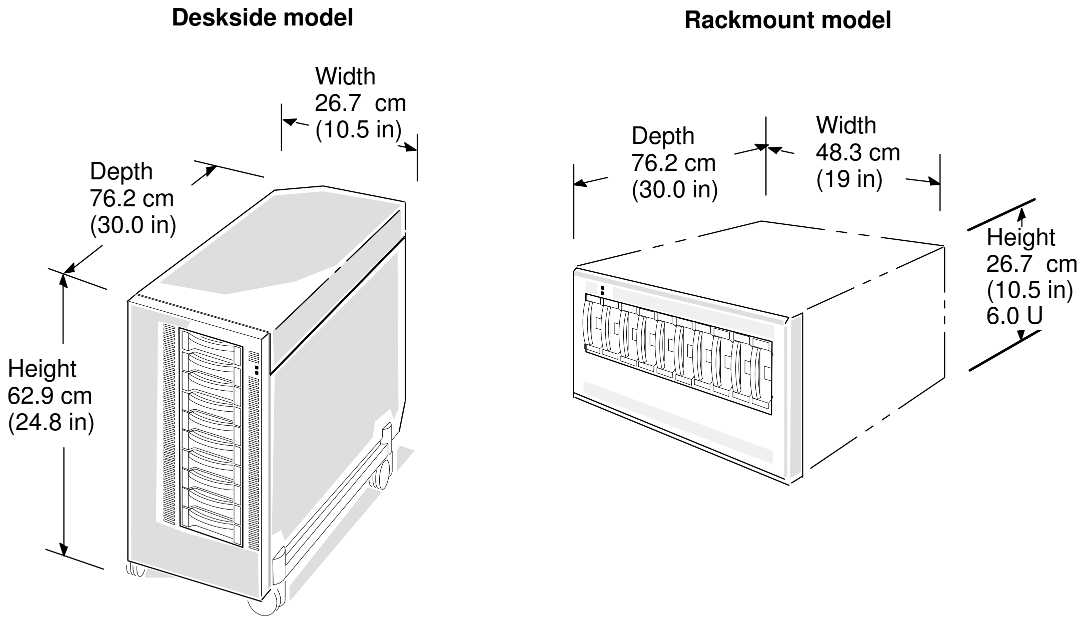

Service clearances	**Deskside**	**Rackmount**
	Front: 81.3 cm (32 in) Rear: same	Front: 81.3 cm (32 in) Rear: same

Weight (without packaging)	**Deskside**	**Rackmount**
Maximum (10 disk modules, 2 SPs, 2 VSCs, BBU)	70.2 kg (156 lb)	63 kg (140 lb)

Operating environment	
Temperature	10ºC to 38ºC (50º F to 100º F)
Relative humidity (noncondensing)	20% to 80%
Altitude	2438 m (8000 ft)
Heat dissipation (max)	2070x10^3 Joules, 1963 BTU/hour

Power requirements	
Voltage rating	90 V ac to 264 V ac, single-phase, 47 to 63 Hz; power supplies are auto-sensing and auto-ranging
Power consumption	600 VA max (apparent), 575 VA max (true)
Phase	1 (USA), or 2 or 3 (Outside USA)

Power cables	
USA primary power	1.8 m (6.0 ft); NEMA L5-15P connection, L5-15 mating connection
Other locations	Local standard ac connection

7-slot storage system requirements

Deskside model

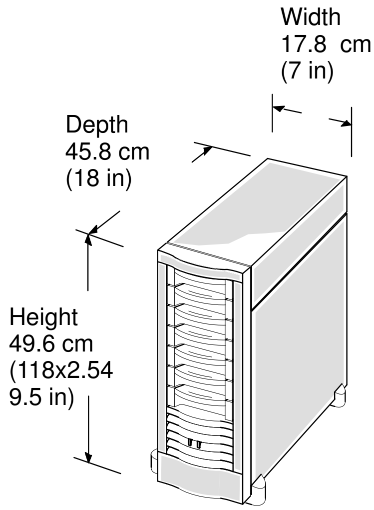

Service clearances	**Deskside**
	Front: 81.3 cm (32 in) Rear: same

Weight (without packaging)	**Deskside**
Maximum (7 disk modules, 1 SP, 2 VSCs)	29.5 kg (65 l5 lb)

Operating environment	
Temperature	10^oC to 38^oC (50^o F to 100^o F)
Relative humidity (noncondensing)	20% to 80%
Altitude	2438 m (8000 ft)
Heat dissipation (max)	$1044x10^3$ Joules, 912 BTU/hour max

Power requirements	
Voltage rating	90 V ac to 264 V ac, single-phase, 47 to 63 Hz; power supplies are auto-sensing and auto-ranging
Power consumption	300 VA max (apparent), 290 VA max (true – typically less than 250 VA)
Phase	1 (USA), or 2 or 3 (Outside USA)

Power cables	
USA primary power	1.8 m (6.0 ft); NEMA L5-15P connection, L5-15 mating connection
Other locations	Local standard ac connection

A note on SCSI cabling

For any but the simplest installation, you should take into account cable lengths to make sure that the total length of any SCSI bus does not exceed the limit of 25 m (81 feet). Different SCSI cables are available; for further details, please refer to your Bull representative.

The two figures below depict sample configurations that make use of different SCSI cabling elements.

Sample configuration 1:

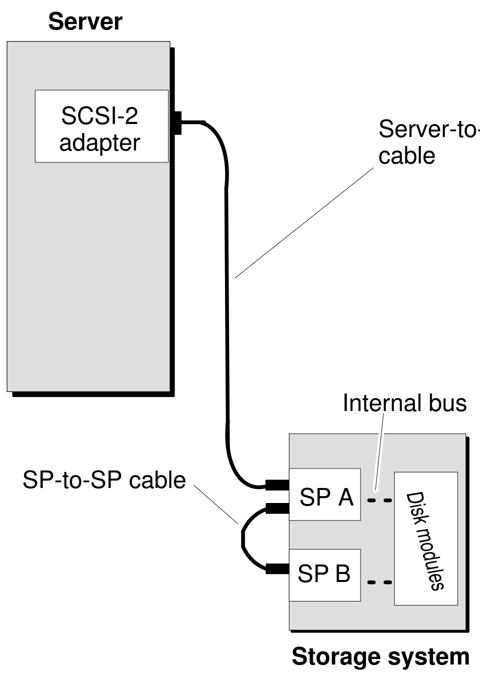

Sample configuration 2:

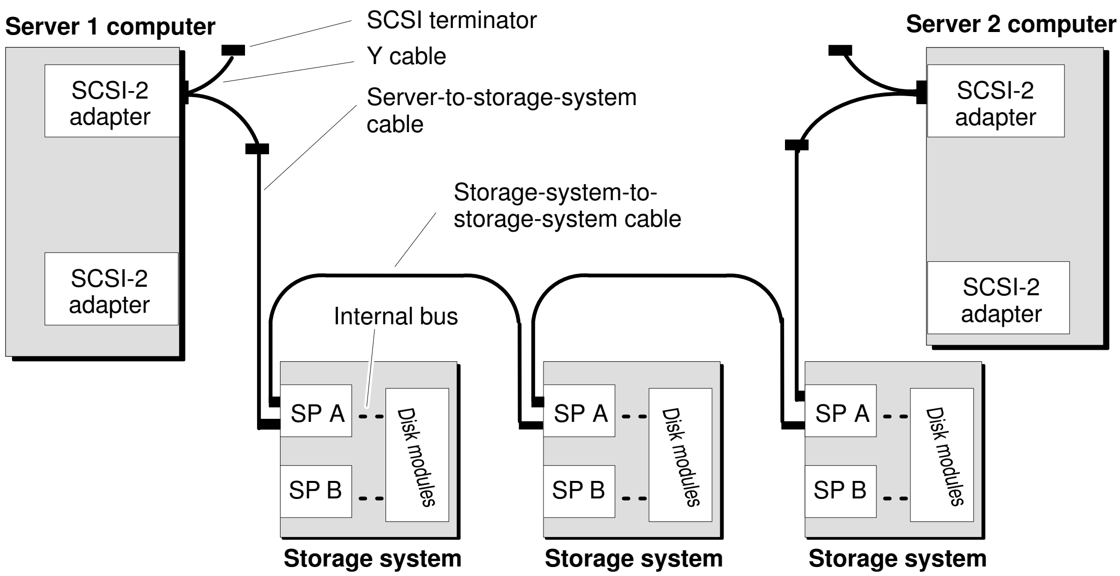

Note: For the second SCSI bus (since the figure above is the basis for a dual-initiator/dual-bus configuration), the cables would be the same as for the first bus, above.

Configuration planning worksheets

This section explains how to complete a planning worksheet that will help you configure the storage system.

A preliminary worksheet lets you note the file systems you want and the physical disks that will hold them. This worksheet shows

- The internal code database disks. With 30–slot systems, there are five database disks; with 20-slot and 10-slot systems, there are four; with a 7-slot system, there are three. For redundancy, we strongly recommend that you use at least two database disks when you run a storage system.
- The cache vault disks, used to store cache data. On a 30-slot, or 10-slot system, all five vault disks must be present for storage-system caching to work; on a 7-slot system, all three vault disks must be present.

For a 30-slot storage system, use the worksheet on page 47 or 48; for a 20-slot storage system, use the worksheet on page 49; for a 10-slot system, use the worksheet on page 50; for a 7-slot system, use the worksheet on page 51.

After planning your file systems, you should complete the storage-system disk worksheet on page 52. You will need this information later when you configure the storage system.

Sample file system and storage-system disk planning worksheets appear on pages 59 and 60.

> **IMPORTANT** Generally, you should plan RAID-1/0 groups, RAID-5 groups, RAID-3 groups, and RAID-0 groups first, then RAID-1 mirrored pairs and other mirrors (provided through AIX software mirroring function), and finally individual units and hot spares. This way, you will have good choices available for disk modules in these RAID groups and mirrors.
>
> Do not plan to bind modules of different capacities in a multiple-module disk unit (RAID group or mirror). If you mix capacities, each module bound will be limited to the capacity of the smallest disk module. For example, if you bind a RAID-1 mirrored pair with a 9- and a 4-Gbyte disk module, the resulting physical disk will hold 4 Gbytes; this wastes 5 Gbytes of space on the larger module.

File System Planning Worksheet, 30-Slot Deskside Storage System

File systems

Drive/file system/volume and disk configuration

1 ____________________

____

2 ____________________

____

3 ____________________

____

4 ____________________

____

5 ____________________

____

File systems

Drive/file system/volume and disk configuration

6 ____________________

____

7 ____________________

____

8 ____________________

____

9 ____________________

____

10 ____________________

____

Disk modules

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

File System Planning Worksheet, 30-Slot Rackmount Storage System

File systems

Drive/file system/volume and disk configuration

1 ____________________

____

2 ____________________

____

3 ____________________

____

4 ____________________

____

5 ____________________

____

File systems

Drive/file system/volume and disk configuration

6 ____________________

____

7 ____________________

____

8 ____________________

____

9 ____________________

____

10 ____________________

____

Disk modules

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

File System Planning Worksheet, 20-Slot Storage System

Drive/file system/volume and disk configuration

1

2

3

4

5

6

Disk modules deskside

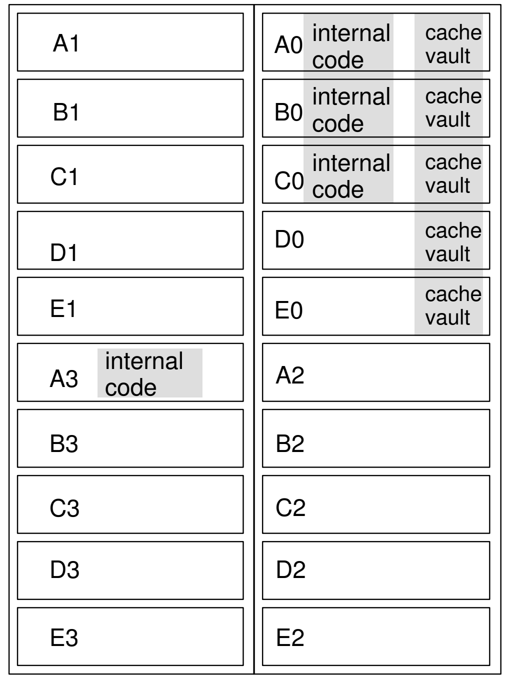

Disk modules, rackmount

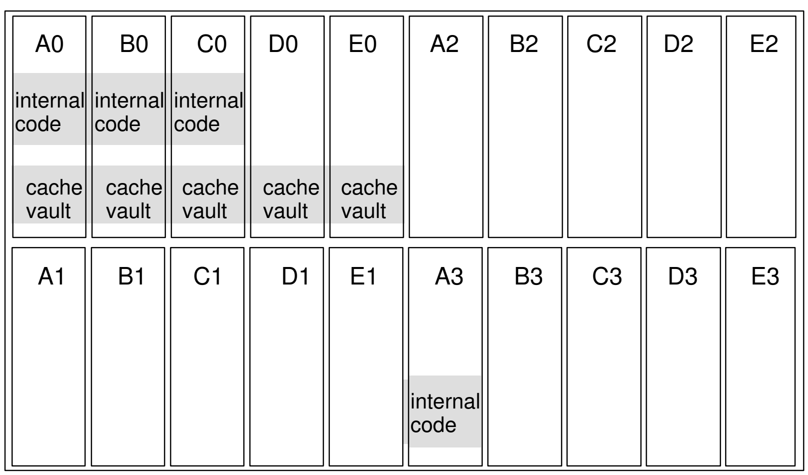

File System Planning Worksheet, 10-Slot Storage System

Drive/file system/volume and disk configuration

1

____

____

2

____

____

3

____

____

4

____

____

5

____

____

6

____

____

Disk modules, Deskside 10-slot

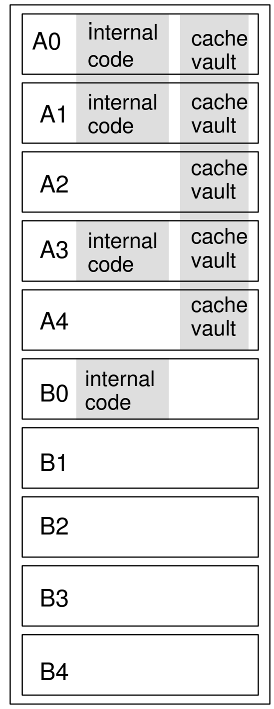

Disk modules rackmount

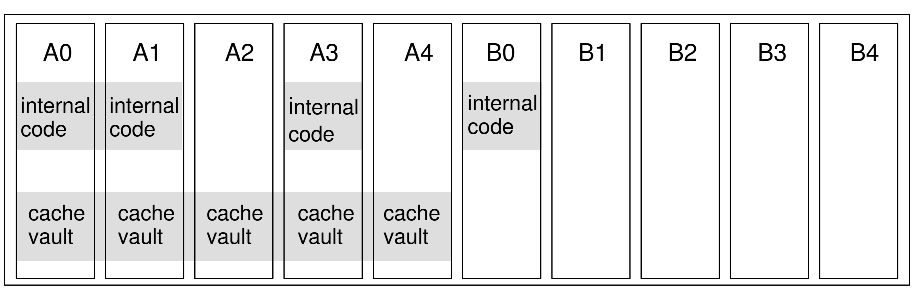

File System Planning Worksheet, 7-Slot Storage System

Drive/file system/volume and disk configuration

1 ______________________________________

____

2 ______________________________________

____

3 ______________________________________

____

4 ______________________________________

____

Disk modules
7-slot

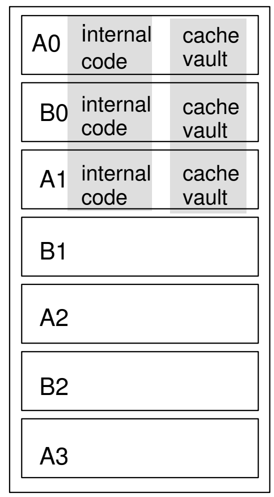

Storage-System Disk (LUN) Planning Worksheet

Server hostname(s):___________________________

Storage-system configuration

☐ Basic ☐ Dual-SP ☐ Dual-adapter/dual-SP ☐ Split bus ☐ Dual-initiator ☐ Dual-initiator/Dual-bus

Storage-system memory (Mbytes): SP A: _____ SP B: _____

Use memory for read/write caching ☐ Cache page size:____Kbytes

Use memory for RAID–3 ☐

Physical disk unit (LUN) number:_____ SIze, Gbytes:____ Disk module IDs: ___________________SP: ☐A ☐B

Disk configuration: ☐ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size:_____Rebuild time:_____ Caching: Read/write ☐ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________________

Software mirror: ☐ Other image disk module ID(s):_______________Device name:_________________

Drive/file system/volume name:____

____

Physical disk unit (LUN) number:_____ SIze, Gbytes:____ Disk module IDs: ___________________SP: ☐A ☐B

Disk configuration: ☐ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size:_____Rebuild time:_____ Caching: Read/write ☐ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________________

Software mirror: ☐ Other image disk module ID(s):_______________Device name:_________________

Drive/file system/volume name:____

____

Physical disk unit (LUN) number:_____ SIze, Gbytes:____ Disk module IDs: ___________________SP: ☐A ☐B

Disk configuration: ☐ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size:_____Rebuild time:_____ Caching: Read/write ☐ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________________

Software mirror: ☐ Other image disk module ID(s):_______________Device name:_________________

Drive/file system/volume name:____

____

Physical disk unit (LUN) number:_____ SIze, Gbytes:____ Disk module IDs: ___________________SP: ☐A ☐B

Disk configuration: ☐ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size:_____Rebuild time:_____ Caching: Read/write ☐ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________________

Software mirror: ☐ Other image disk module ID(s):_______________Device name:_________________

Drive/file system/volume name:____

____

Completing the file system worksheet

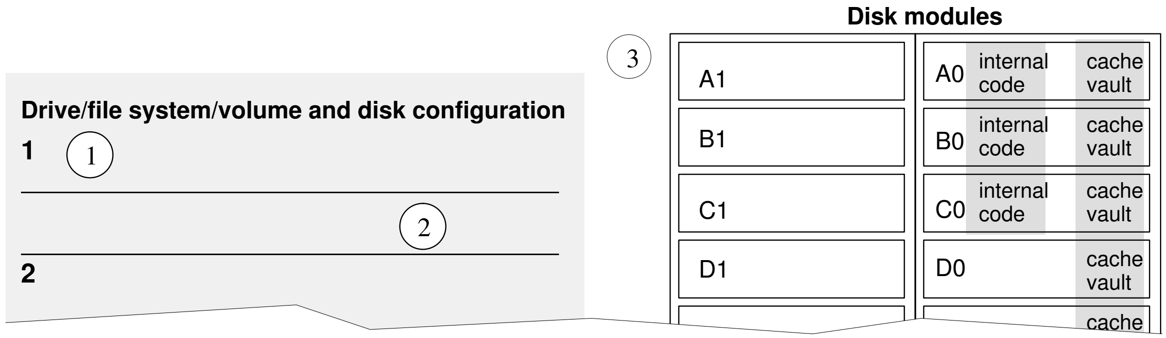

Write information on the file system worksheet as explained below. A sample completed worksheet follows later on page 59.

1 For each **drive, file system** or **logical volume**, write the drive letter (if applicable), the name, and optionally the type of application you plan to run on this disk unit. For example, “S (system disk).” If the disk will be a hot spare and will thus not have a file system, write hot spare in this column.

If you want more than one partition or file system on one physical disk, just repeat the physical disk configuration as often as needed. For example, partition ACCTS, configuration “RAID-5 #1” and file system MAIL, configuration “RAID-5 #1.”

2 The **disk configuration** is the type of unit into which you will bind this disk module. You can choose RAID-5, RAID-3, RAID-1/0, RAID-1, RAID-0, individual unit, or hot spare. For a hot spare, write hot spare in both columns.

For clarity, number each physical disk configuration within its type. For example, with two RAID-5 groups, “S, RAID-5 #1” and “T, RAID-5 #2.”

3 For each disk confguration you defined in entry 2, decide which **disk modules** you want to form each physical disk unit (LUN), circle these disk modules, and write the configuration (for example, RAID-5) in the circle. This is information you will use to bind the disk. For storage-system caching to occur, there must be disk modules in all the cache vault slots.

Completing the storage-system disk (LUN) worksheet

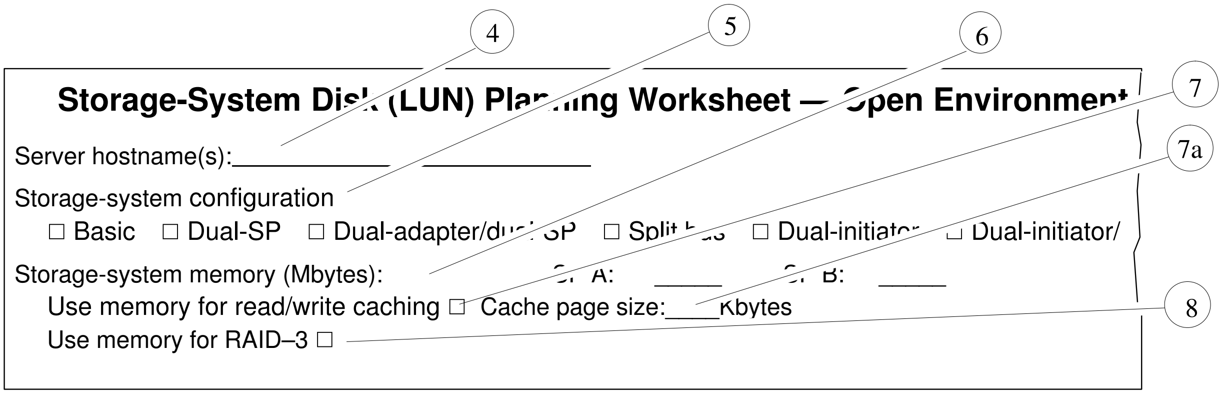

Complete the header portion of the worksheet for each storage system as described below. Copy the blank worksheet as needed. A sample completed worksheet follows later on page 60.

4

For **server hostname(s)**, enter the hostname of the computer these disk units will serve. If this storage system will be used by two servers, write the hostnames of both servers. We suggest you copy this worksheet and complete the copy for the other server, with the primary and secondary names reversed. This is particularly important where one server may take over the other's disks, that is, in a dual-initiator or dual-initiator/dual-bus configuration.

For each storage-system disk whose primary owner will be the other server, write "secondary" in the section "Operating system/file system information."

5

Specify the **storage-system configuration**: basic, dual-SP, dual-adapter/dual-SP, split-bus, dual-initiator, or dual-initiator/dual-bus.

6

Storage-system 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 read/write caching or RAID–3 use.

7

Use memory for read/write caching. You can use the storage system 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, the latest revision of licensed internal code can allocate the memory to read and/or write caches automatically. The benefits of caching are explained starting on page 4; the benefits of RAID–3 are explained on page 26. For caching, check the box and continue; for RAID–3, skip to step 8.

7a

The **cache page size** 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.

8 **Use memory for RAID–3**. If you want to use RAID–3, check the box.

Physical disk unit (LUN) entry

Complete a physical disk unit (LUN) entry for each physical disk and hot spare.

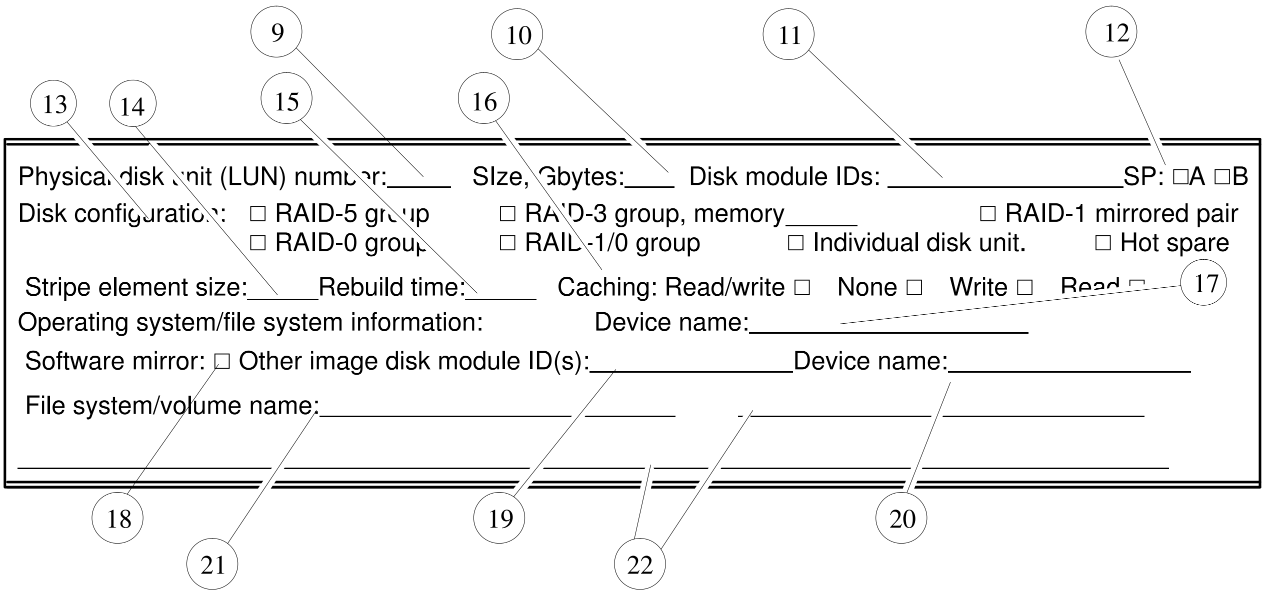

9 The **physical disk unit (LUN) number** is a hexadecimal number assigned when you bind the disk. (The LUN is a logical concept, but the operating system sees the disk as a single physical disk unit, hence the seemingly contradictory names.) By default, the number of the first LUN bound is 0, the second 1, and so on. For example, if the first LUN you bind is a RAID-5 group on disk modules A0 through A4, each module becomes part of LUN 0.

Some specific SCSI adapters allow only eight physical disks (numbers 0 through 7). When such a restriction does apply (depending on the SCSI adapter type), if you want a hot spare, assign the hot spare a unit number above 7; for example, 8 or 9. The operating system never accesses a hot spare, so the unit number is irrelevant to it.

10 For **size**, enter the user-available capacity in gigabytes (Gbytes) of the whole physical disk (LUN). You can determine the capacity as follows:

RAID-5 or RAID-3 group:	module-size X (number-of-modules –1)
RAID 1/0 or RAID-1 group:	(module-size X number-of-modules) /2
RAID-0 group:	module-size X number-of-modules
Individual unit:	module-size

For example,

- A five-disk RAID-5 or RAID-3 group of 9-Gbyte disks holds 36 Gbytes;
- An eight-disk RAID-1/0 group of 9-Gbyte disks also holds 36 Gbytes;
- A RAID-1 mirrored pair of 9-Gbyte disks holds 9 Gbytes; and
- An individual disk of a 9-Gbyte disk also holds 9 Gbytes.

As mentioned earlier, each disk module in the physical disk must have the same capacity; otherwise, you will waste disk storage space.

11 For **disk module IDs**, enter the disk module ID(s) of all modules that will make up the physical disk. For example, for a RAID-5 group with modules A0 through E0, enter A0, B0, C0, D0, and E0.

12 For **SP,** if the storage system has two SPs, specify the SP that will run the unit: SP A or SP B.

13 For **disk configuration**, specify the disk configuration from the first worksheet, entry 2. For example, RAID-5 or hot spare. For a hot spare (not strictly speaking a physical disk unit at all), skip the rest of this physical disk entry and continue to the next physical disk unit 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 6 Mbytes of memory.

14 The **stripe element size** is the size in disk sectors (512 bytes) that the storage system can read or write to a single disk module without requiring access to another module (assuming that the transfer starts at the first sector in the stripe). The size you choose can affect the performance of a RAID-5, RAID-1/0, or RAID-0 group. For a RAID-3 group, the stripe element size is fixed at one sector, a value you cannot change.

The smaller the stripe element size, the more efficient the distribution of data read or written. However, if the stripe size is too small for a single I/O operation, the operation requires access to two stripes, thus causing the hardware to read and/or write from two disk modules instead of one. Generally, we recommend the smallest stripe element size that will rarely force access to another stripe. The default stripe element size is 128 sectors. The size should be an even multiple of 16 sectors (8 Kbytes).

The stripe element size value becomes an integral part of the physical disk unit; you cannot change it without unbinding the disk (eliminating all data on it) and then rebinding.

15 The **rebuild time** is the amount of time that the storage system allots to reconstruct the data on a hot spare or a replacement for a failed disk module. The rebuild time applies to all RAID groups except RAID-0. The time you specify determines the amount of resources the SP devotes to rebuilding instead of normal I/O activity. The default time of 4 hours is adequate for most situations. A rebuild time of 2 hours rebuilds the disk quicker, but degrades response time slightly. A rebuild time of 0 hours rebuilds as quickly as possible but degrades performance significantly. If your site requires fast response time and you want to minimize degradation to normal I/O activity, you can extend the rebuilding process over a longer period of time, such as 24 hours. You can change the rebuild time after binding without damaging the information stored on the physical disk unit. Enter the integer number of hours you want, or 0.

16 For **caching**, specify whether you want caching for this unit. You can also specify read or write caching. (If you will use automatic cache allocation, the software will select both kinds.) The ability to specify caching on a physical unit 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 disk configurations

RAID-5	RAID-3	RAID-1	RAID-1/0	RAID-0	Individual unit
Highly Recommended	Not allowed	Recommended	Recommended	Recommended	Recommended

Operating system/file system information

17 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.

18 A **software mirror** is useful for making individual units and RAID-0 groups highly available. If you plan to use AIX software mirroring with this physical disk, check here. If not, skip to entry 21.

19 For the **other image disk module ID(s)**, enter the disk module ID(s) of all modules that make up the other image's physical disk. For highest availability, specify disks on different internal buses (A, B, C, D, E) than those for the primary image's modules.

20 For the secondary image, 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 now.

21 For the **file system/volume name**, write the name of the file system or logical volume you will create on this physical disk unit. This is the same name you wrote on the file system worksheet (item 1).

22 On these lines, write any pertinent notes; for example, the file system mount- or graft-point directory pathname (from the root directory). If this storage system is to be shared with another server, and the other server is the primary owner of this disk, write "secondary." (As mentioned in entry 4, if the storage system will be used by two servers, we suggest you complete one of these worksheets for each server.)

Sample worksheets

The sample worksheets in this section show a single-server configuration: a dual-adapter-dual-SP configuration. The disk-array storage system will have two SCSI-2 adapters, two SPs, and three different kinds of physical disk:

- RAID-5 group,
- RAID-1 mirrored pair,
- individual disk unit.

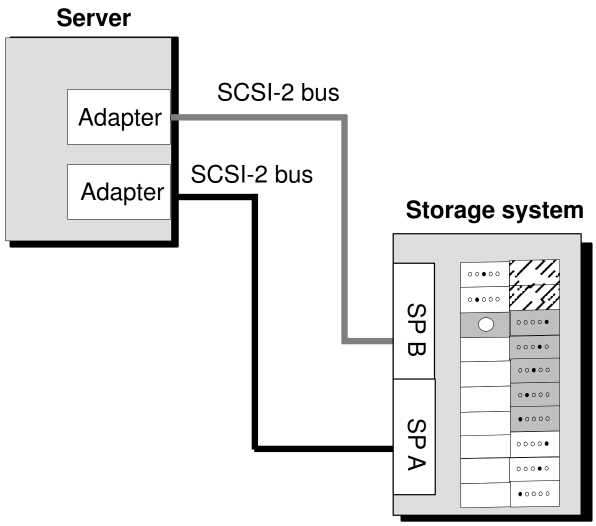

If the disk modules shown here hold 9-Gbytes, then the storage system will provide 36 Gbytes of disk storage with 32 Gbytes highly available. The disk configurations are as follows:

Unit **S** on two disk modules bound as a RAID-1 mirrored pair for 9 Gbytes of storage; user directories.

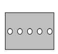

Unit **T** on five disk modules bound as a RAID-5 group (individual access array) for 36 Gbytes of storage; for the database.

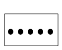

Unit **U** on five modules bound as a RAID-5 group for 36 Gbytes of storage; for clients and mail.

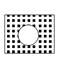

Unit **V** on one disk module bound as an individual unit for 9 Gbytes of storage; temporary directory for users.

Page 59 shows the file system planning worksheet, and page 60 shows the physical disk planning worksheet, for the sample system.

File System Planning Worksheet, 20-Slot Storage System

Drive/file system/volume and disk configuration

1 S (user dirs,
Volume label RAID1_S) RAID·1 #1

2 T (database, Volume
label RAID5_T) RAID·5 #1

3 U (clients, mail, Volume
label RAID5_U) RAID·5 # 2

4 V (temp Volume label DISK_V)
indiV unit #1

5

6

Disk modules deskside

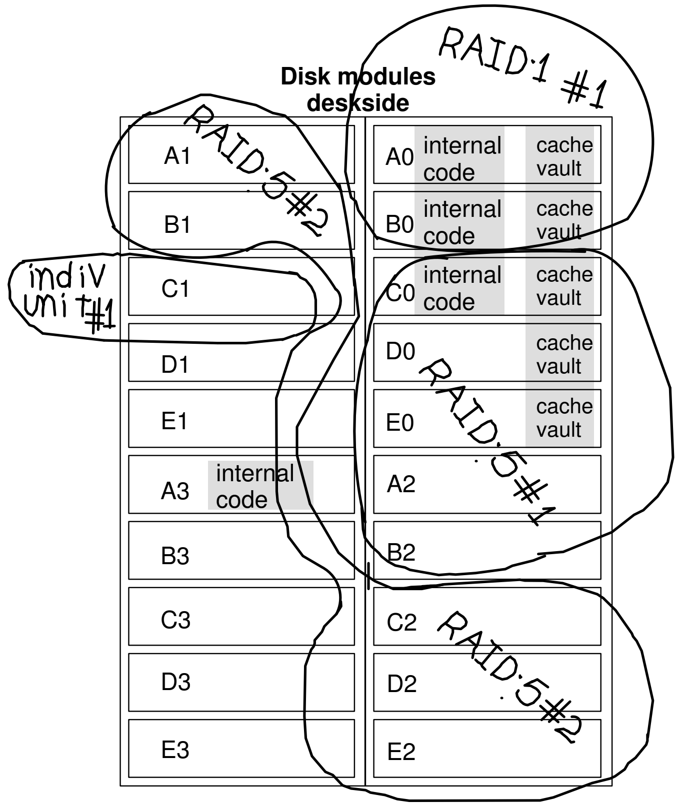

Disk modules, rackmount

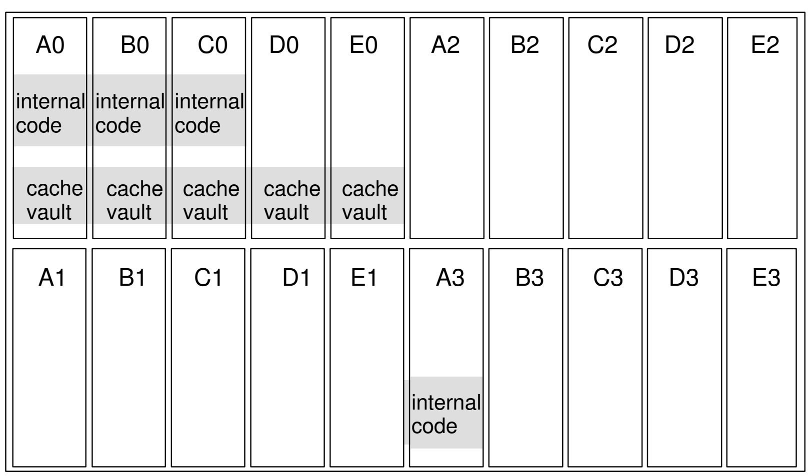

Storage-System Disk (LUN) Planning Worksheet — Open Environments

Server hostname(s): accounts_rec

Storage-system configuration

☐ Basic ☐ Dual-SP ☒ Dual-adapter/dual-SP ☐ Split bus ☐ Dual-initiator ☐ Dual-initiator/Dual-bus

Storage-system memory (Mbytes): SP A: 32 SP B: 32

Use memory for read/write caching ☒ Cache page size: 2 Kbytes

Use memory for RAID–3 ☐

Physical disk unit (LUN) number: 0 SIze, Gbytes: 9 Disk module IDs: A0,B0 SP: ☒A ☐B

Disk configuration: ☐ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☒ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size: n/a Rebuild time: 2 hrs Caching: Read/write ☒ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________

Software mirror: ☐ Other image disk module ID(s):__________Device name:__________

Drive/file system/volume name: S user dirs, Volume label RAID1_S

Physical disk unit (LUN) number:_____ SIze, Gbytes: 36 Disk module IDs: C0, D0, E0, A2, B2 SP: ☒A ☐B

Disk configuration: ☒ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size: 128 Rebuild time: 1 hr Caching: Read/write ☒ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________

Software mirror: ☐ Other image disk module ID(s):__________Device name:__________

Drive/file system/volume name: T (database, Volume label RAID5_T)

Physical disk unit (LUN) number: 2 SIze, Gbytes: 36 Disk module IDs: C2, D2, E2, A1, B1 SP: ☐A ☒B

Disk configuration: ☒ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☐ Individual disk unit ☐ Hot spare

Stripe element size: 128 Rebuild time: 2 hrs Caching: Read/write ☒ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________

Software mirror: ☐ Other image disk module ID(s):__________Device name:__________

Drive/file system/volume name: U (clients, mail, Volume label RAID5_U)

Physical disk unit (LUN) number: 3 SIze, Gbytes: 9 Disk module IDs: C1 SP: ☐A ☒B

Disk configuration: ☐ RAID-5 group ☐ RAID-3 group, memory_____Mbytes ☐ RAID-1 mirrored pair

☐ RAID-0 group ☐ RAID-1/0 group ☒ Individual disk unit ☐ Hot spare

Stripe element size: n/a Rebuild time: n/a Caching: Read/write ☐ None ☐ Write ☐ Read ☐

Operating system/file system information: Device name:____________

Software mirror: ☐ Other image disk module ID(s): n/a Device name: n/a

Drive/file system/volume name: V temporary user directory, Volume label DISK_V

End of Document

Index

D

E

F

G

H

S

Vos remarques sur ce document / Technical publication remark form

Titre / Title : Planning a DAS Installation (SCSI Environments)

Nº Reférence / Reference Nº : 86 A1 84GX 01

Daté / Dated : June 1998

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 ELECTRONICS EUROPE S.A.
Service CEDOC
331 Avenue PATTON – BP 428
49004 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 ELECTRONICS EUROPE S.A.
Service CEDOC
ATTN / MME DUMOULIN
331 Avenue PATTON – BP 428
49004 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 # Nº Référence CEDOC	**Qty** Qté	**CEDOC Reference #** Nº Référence CEDOC	**Qty** Qté	**CEDOC Reference #** Nº 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.

PLACE BAR CODE IN LOWER
LEFT CORNER

BULL ELECTRONICS EUROPE S.A.
Service CEDOC
331 Avenue PATTON – BP 428
49004 ANGERS CEDEX 01
FRANCE

ORDER REFERENCE
86 A1 84GX 01

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

DAS

Planning a
DAS Installation
(SCSI Environments)

86 A1 84GX 01

DAS

Planning a
DAS Installation
(SCSI Environments)

86 A1 84GX 01

DAS

Planning a
DAS Installation
(SCSI Environments)

86 A1 84GX 01