

# StorageWorks™ Solutions

---

## Configuration Planning Guide

Order Number: EK-BA350-PG. A01

This Guide describes planning the hardware for a StorageWorks storage system configuration. It uses, as a case study, a fictional enterprise with varied storage requirements to illustrate the decision making processes that are part of configuration. The Guide lists the rules for configuring a hardware system, and provides suggestions to help ensure the configuration proceeds smoothly from the initial requirements' analysis to the final implementation on the customer site.

---

**September 1995**

While Digital believes the information included in this publication is correct as of the date of publication, it is subject to change without notice.

Digital Equipment Corporation makes no representations that the interconnection of its products in the manner described in this document will not infringe existing or future patent rights, nor do the descriptions contained in this document imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

**NOTE:** This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications.

**FCC ID: A09-BA350****Date of Original Grant: January 11, 1994****Date of Class 2 Change Authorization: April 26, 1994**

**NOTE:** This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation.

Any changes or modifications made to this equipment may void the user's authority to operate this equipment.

This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If this equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna
- Increase the separation between the equipment and receiver
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected
- Consult the dealer or an experienced radio/TV technician for help

© Digital Equipment Corporation 1992, 1993, 1994.

Printed in the United States of America.

All rights reserved.

applicationDEC, AXP, CI, DEC, DECpc, DECraid, DECstation, DECsystem, Digital, HSC, HSC50, HSC95, HSJ, HSJ40, HSR95, InfoServer, MicroVAX, MicroVAX 3100, OpenVMS, RX, RZ, StorageWorks, TZ, ULTRIX, VAX, VAX 4000, VAXstation, VAXstation 4000, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

Hewlett-Packard is a registered trademark of HewlettPackard Company.

OSF and OSF/1 are trademarks of the Open Software Foundation, Inc.

Sun Microsystems is a registered trademark of Sun Microsystems, Inc.

UNIX is a registered trademark of X/Open Co. Ltd., Inc.

All other trademarks and registered trademarks are the property of their respective holders.

The postpaid READER'S COMMENTS card requests the user's critical evaluation to assist in preparing future documentation.

This document was prepared using MICROSOFT WORD Version 6.0.

---

# Contents

<b>Preface .....</b>	<b>vii</b>
Documentation Conventions .....	viii
<b>1 Introduction</b>	
<b>2 Analyzing the Storage Requirements</b>	
2.1 Analyzing Access to the Storage System .....	2-6
2.1.1 Determine the User Groups That Will Access the Storage System .....	2-6
2.1.2 Determine the Applications That Will Access the Storage System? .....	2-6
2.2 Determine the Required Storage Capacity .....	2-6
2.3 Identify the Data Categories .....	2-7
2.4 Determine Each Application's I/O Requirements .....	2-8
2.4.1 I/O Activity.....	2-8
2.4.2 Data Throughput .....	2-8
2.4.3 I/O Response Time .....	2-9
2.5 Determine the Required Level of Data Availability.....	2-9
2.6 Summarizing the Analysis .....	2-10
<b>3 Planning the StorageWorks Configuration</b>	
3.1 Planning for Storage of Active Data.....	3-5
3.1.1 Planning for Capacity .....	3-5
3.1.2 Planning for Storage System Performance.....	3-7
3.1.2.1 Controller Caching.....	3-7
3.1.2.2 Multiple Independent Controllers .....	3-10
3.1.2.3 Using 8 or 16-Bit SBB Shelves.....	3-12
3.1.3 Planning for Data Availability .....	3-13
3.1.3.1 Using Redundant Controllers.....	3-13
3.1.3.2 Using Redundant Power Supplies .....	3-13
3.1.4 Using RAID to Improve Performance and Availability.....	3-16
3.1.5 Tradeoffs Between Performance and Availability.....	3-19
3.2 Planning for Storage of Dormant and Inactive Data .....	3-19
3.3 Planning for Backup and Restore Operations .....	3-19
<b>4 Selecting StorageWorks Products</b>	
4.1 Selecting Controllers and Adapters .....	4-2
4.2 Selecting Storage Devices .....	4-5

4.2.1	Selecting the Device Types .....	4-5
4.2.2	Determining the Number of Storage Devices.....	4-6
4.3	Selecting Shelves .....	4-9
4.3.1	Device Shelves.....	4-10
4.3.2	Controller Shelves .....	4-14
4.3.3	Combination Shelves .....	4-15
4.4	Selecting Power Components .....	4-16
4.5	Selecting Enclosures .....	4-16
4.6	Selecting Accessories .....	4-17
4.7	The Configuration Planned For the Fictional Enterprise.....	4-17
<b>5</b>	<b>Implementing the StorageWorks Configuration</b>	
5.1	Ordering the StorageWorks Components.....	5-1
5.2	Preparing the Site.....	5-1
5.3	Installing the Equipment.....	5-2
5.4	Integrating the Storage System.....	5-2
5.5	Acceptance Testing of the Storage System.....	5-3
5.6	Populating the Storage System With Data .....	5-4
<b>A</b>	<b>General Configuration Rules</b>	
A.1	Controller Configuration Rules.....	A-1
A.2	SBB Shelf Configuration Rules .....	A-1
A.2.1	3.5-Inch SBB Configurations on 3-Port Controllers.....	A-3
A.2.2	5.25-Inch SBB Configurations on 3-Port Controllers.....	A-3
A.2.3	3.5-Inch SBB Configurations on 6-Port Controllers.....	A-4
A.2.4	5.25-Inch SBB Configurations on 6-Port Controllers.....	A-4
A.3	Power System Configuration Rules.....	A-5
A.4	SCSI Bus Configuration Rules.....	A-5
A.4.1	Terminating SCSI Buses .....	A-6
A.4.2	Computing SCSI Bus Lengths.....	A-7
A.4.3	Single StorageWorks Shelf SCSI Bus Lengths .....	A-7
A.4.4	Multiple StorageWorks Shelves SCSI Bus Lengths .....	A-10
<b>B</b>	<b>RAID Overview</b>	
B.1	Striping (RAID Level 0) .....	B-2
B.2	Mirroring (RAID Level 1) .....	B-3
B.3	RAID Level 3 .....	B-4
B.4	Parity RAID (RAID Level 5) .....	B-5
B.5	Striped Mirrorsets (RAID Level 0 & 1).....	B-6
B.6	Summary Comparison of RAID Levels.....	B-7
<b>C</b>	<b>Migration Planning — An Overview</b>	
C.1	Identify the Current System In Terms of Its I/O Requirements and Performance.....	C-1
C.1.1	Identify Any System Level or Storage Subsystem Performance Problems.....	C-2
C.1.2	Identify the Applications, Software, and Uses of the Current System .....	C-2
C.1.3	Determine the Expected Storage Capacity and Performance Requirements .....	C-2

C.1.4 Identify Which Part of the Storage System Will Be Migrated .....	C-2
C.1.5 Identify the Data Backup Methodologies and Products Used.....	C-2
C.1.6 Identify the System and Storage/Controller Configuration.....	C-3
C.1.7 Identify the Methods Used to Provide Device Availability.....	C-3
C.1.8 Determine if the Disk Storage is Fragmented .....	C-3
C.2 Determine the Types and Numbers of Replacement Storage Products .....	C-4

## StorageWorks Glossary

### Figures

Figure 1–1 Typical StorageWorks Configuration Using Redundant Controllers.....	1-2
Figure 1–2 Typical StorageWorks Configuration Using Independent Controllers.....	1-3
Figure 3–1 Storage Relationships.....	3-5
Figure 3–2 Independent Multiple Controllers Configuration .....	3-11
Figure 3–3 Redundant Controller Configuration .....	3-13
Figure 3–4 Cost–Availability–Performance Triangle .....	3-16
Figure 4–1 Typical Device Shelf .....	4-10
Figure 4–2 Typical Controller Shelf.....	4-14
Figure 4–3 Typical Combination Shelf .....	4-15
Figure 4–4 View of SW800 Enclosure Housing the Fictional Configuration .....	4-18
Figure 4–5 Logical Configuration of the Device Shelves for the CAD and File Services Applications .....	4-19
Figure A–1 SCSI Bus Segments — Single SBB Shelf.....	A-8
Figure A–2 SCSI Bus Segments — Multiple SBB Shelves .....	A-11

### Tables

Table 1 StorageWorks Related Documentation.....	viii
Table 1–1 Organization of the Configuration Guide.....	1-4
Table 2–1 Data Category and Storage Type.....	2-7
Table 2–2 Procedures for Completing the Summary Worksheet.....	2-11
Table 3–1 Configuration Planning Worksheet and Product Selection .....	3-1
Table 4–1 StorageWorks Shelf Capacities .....	4-9
Table 4–2 Device Shelf Requirements for 3.5 Inch Devices.....	4-10
Table 4–3 Device Shelf Requirements for 5.25 Inch Devices.....	4-11
Table 5–1 Site Preparation Worksheet .....	5-2
Table 5–2 Pre-Integration Checklist.....	5-3
Table A–1 3.5-Inch SBB Configurations on 3-Port Controllers .....	A-3
Table A–2 5.25-Inch SBB Configurations on 3-Port Controllers .....	A-3
Table A–3 3.5-Inch SBB Configurations on 6-Port Controllers .....	A-4
Table A–4 5.25-Inch SBB Configurations on 6-Port Controllers .....	A-4
Table A–5 SCSI Bus Parameters.....	A-6
Table A–6 Single SBB Shelf SCSI Bus Lengths.....	A-9
Table A–7 Multiple SBB Shelves SCSI Bus Lengths.....	A-10



---

# Preface

The StorageWorks Solutions Configuration Planning Guide describes planning the hardware for a StorageWorks storage system configuration. It uses, as a case study, a fictional enterprise with varied storage requirements to illustrate the decision making processes that are part of configuration. The Guide lists rules for configuring a system, and provides suggestions to help ensure the configuration proceeds smoothly from the initial analysis to the final installation on the customer site.

## Intended Audience

This manual is for use by personnel responsible for designing and configuring a StorageWorks system. This includes site planners and system managers.

## Structure

This manual is organized as follows:

Chapter	Description
Chapter 1 Introduction	Introduces configuration and describes the remainder of the Configuration Guide.
Chapter 2 Analyzing the Storage Requirements	Discusses analysis of the storage requirements.
Chapter 3 Planning the StorageWorks Configuration	Provides a step by step approach to planning the storage configuration.
Chapter 4 Selecting StorageWorks Products	Provides suggestions for selecting the StorageWorks products required to implement the configuration.
Chapter 5 Implementing the StorageWorks Configuration	Provides suggestions for completing the configuration.
Appendix A General Configuration Rules	Provides general configuration rules for StorageWorks shelves, power systems, and the SCSI bus.
Appendix B RAID Overview	Provides a brief introduction to RAID as supported for StorageWorks products
Appendix C Migration Planning Overview	Provides a high level overview of migration planning

## Documentation Conventions

The following conventions are used in this manual:

<b>boldface type</b>	Boldface type indicates the first instance of terms being defined in text, in the glossary, or both.
<i>italic type</i>	Italic type indicates emphasis and complete manual titles. In the glossary, italic type is also used to indicate cross-references.

## Related Documents

Table 1 lists the StorageWorks-related user documents organized by use, system, or product.

**Table 1 StorageWorks Related Documentation**

<b>Document Title</b>	<b>Order Number</b>
<b>StorageWorks Primary Publications</b>	
<i>StorageWorks Solutions Configuration Guide</i>	EK-BA350-CG
<i>StorageWorks Solutions Shelf and SBB User's Guide</i>	EK-BA350-UG
<b>StorageWorks RAID Array 110 Subsystem</b>	
<i>DEC RAID Utilities User's Guide</i>	EK-DECRA-UG
<i>HSZ10-AA Controller Site Preparation Guide</i>	EK-HSZ10-IN
<i>StorageWorks Expansion Storage Pedestal User's Guide</i>	EK-SMCPA-UG
<i>StorageWorks RAID Array 110 Subsystem User's Guide</i>	EK-SM2CA-UG
<i>StorageWorks RAID Array 110 Utility for MS-DOS User's Guide</i>	AA-Q0N5A-TE
<i>StorageWorks RAID Array 110 Utility for Novell NetWare User's Guide</i>	AA-Q0N4A-TE
<i>StorageWorks RAID Array 110 Utility for SCO UNIX User's Guide</i>	AA-Q0N6A-TE
<b>StorageWorks Array Controller 140-Series</b>	
<i>StorageWorks Array Controller HS Family of Array Controllers User's Guide</i>	EK-HSFAM-UG
<i>StorageWorks BA350-MA Controller Shelf User's Guide</i>	EK-350MA-UG
<b>DECraid™+ Rackmount StorageSubsystem and HSC[™] Intelligent I/O Servers</b>	
<i>HSC Controller User's Guide</i>	AA-PFSQA-TK
<i>HSC Controller Installation Manual</i>	EK-HSCMN-IN

(continued on next page)



**Table 2 (Cont) StorageWorks Related Documentation**

<b>StorageWorks Enclosures</b>	
<i>BA35X-VA Vertical Mounting Kit User's Guide</i>	EK-350SV-UG
<i>InfoServer 1000 Installation and Owner's Guide</i>	EK-INFLC-OM
<i>StorageWorks Solutions Desktop Expansion Unit User's Guide</i>	EK-BA353-UG
<i>StorageWorks Solutions SW500 and SW800 Cabinet Metric Shelf Bracket Kit Installation Guide</i>	EK-35XRD-IG
<i>StorageWorks Solutions HSR95-Series Cabinet Metric Shelf Bracket Kit Installation Guide</i>	EK-35XRC-IG
<i>StorageWorks RETMA Shelf Rail Kit Installation Guide</i>	EK-35XRB-IG
<i>StorageWorks Solutions SW500-Series Cabinet Installation and User's Guide</i>	EK-SW500-IG
<i>StorageWorks Solutions SW500-Series Cabinet Cable Distribution Unit Installation Guide</i>	EK-SW5CU-IG
<i>StorageWorks Solutions SW800-Series Data Center Cabinet Cable Distribution Unit Installation Guide</i>	EK-SWCUD-IS
<i>StorageWorks Solutions SW800-Series Data Center Cabinet Installation and User's Guide</i>	EK-SW800-IG
<i>Instructions for Installation of the Tape Bezel Kit on SF400 Series Cabinets*</i>	EK-SFTBK-IS
<b>Storage Devices SCSI</b>	
<i>Signal Converter DWZZA-AA</i>	EK-DWZZA-SV
<i>SCSI Signal Converter DWZZA-MA</i>	EK-DWZZM-SV
<i>SCSI Signal Converter DWZZA-VA</i>	EK-DWZAA-SV
<i>StorageWorks Building Blocks User's Guide</i>	EK-SBB35-UG
<i>StorageWorks Solutions 3½-Inch Storage Device Installation Guide</i>	EK-MC350-IG
<i>StorageWorks Solutions 5¼-Inch Storage Device Installation Guide</i>	EK-MC525-IG
<i>TZ86-VA Cartridge Tape Drive Device Addresses</i>	EK-TZ86A-UG
<i>TZ87-VA Cartridge Tape Drive Device Addresses</i>	EK-TZ87A-DA
<i>StorageWorks Solutions TZ8x7-Series Tape Drive SCSI Bus Configuration and Installation Guide</i>	EK-TZ8X7-IG
<b>General Reference Publications</b>	
<i>Digital Systems and Options Catalog</i>	†
<i>StorageWorks Solutions Products Catalog</i>	EK-BA350-PC
<i>Small Computer System Interface, An Overview</i>	EK-SCSIS-OV
<i>Small Computer System Interface, A Developer's Guide</i>	EK-SCSIS-DK
* Applies to StorageWorks cabinets	
† Available from your Digital account representative	



---

# Introduction

---

## Note

---

Storage configuration requires both hardware and software configurations. While the two are interdependent, software configuration is beyond the scope of this guide. Consequently, storage configuration in this guide refers to hardware configuration.

---

Storage configuration is the process of creating a storage system. Since both users and applications require some amount of storage for their exclusive use, the storage system must include the storage requirements for both. These storage requirements are defined in terms of:

- Storage capacity
- I/O requirements
- Data availability
- Cost constraints

Because applications and users have varied storage requirements, the storage system can consist of multiple configurations; each configuration designed to meet a specific set of requirements. However, the combined configurations provide a single StorageWorks solution.

This guide helps you provide a StorageWorks solution by guiding you through analyzing the storage requirements, planning a solution, selecting the StorageWorks products, and implementing the storage system.

During the analysis phase you define the purpose of the host system, then determine the storage requirements in terms of storage capacity, I/O requirements, and data availability.

During the planning phase you design a StorageWorks solution to satisfy the storage capacity requirements and meet the goals specified by the I/O requirements and data availability. The planning phase generally requires that you balance your requirements for I/O, data availability, and capacity to achieve the best compromise.

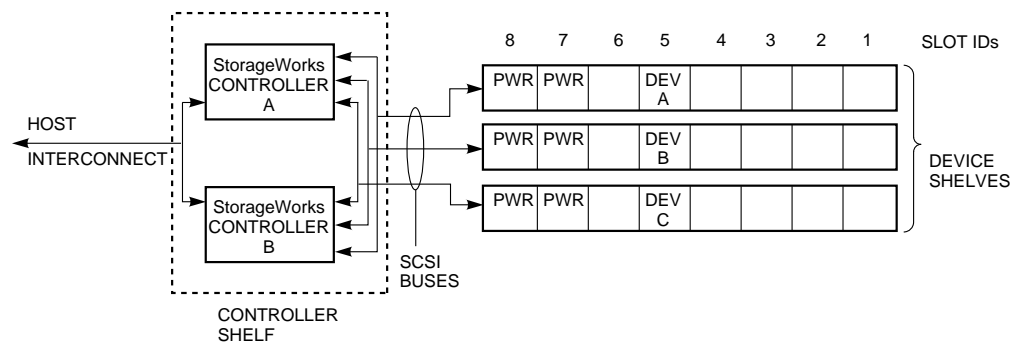
During the device selection phase you choose the StorageWorks products that, when combined into a storage system, meet or exceed the requirements defined during the analysis phase while meeting the cost constraints. This may include selecting any or all of the following StorageWorks products:

- Controllers and adapters
- Storage devices
- Shelves
- Power components
- Enclosures
- Cables and accessories

Figures 1–1 and 1–2 show typical configurations using some of the products from the previous list. Figure 1–1 shows a configuration using redundant controllers mounted in one controller shelf and three storage devices mounted in three devices shelves. The device shelves connect to the three SCSI ports of both controllers. Figure 1–2 shows a configuration using multiple independent controllers, each in a separate controller shelf, and six device shelves. Each device shelf contains two storage devices. Three device shelves connect to the three SCSI ports of controller A, the remaining shelves connect to the three SCSI ports of controller B.

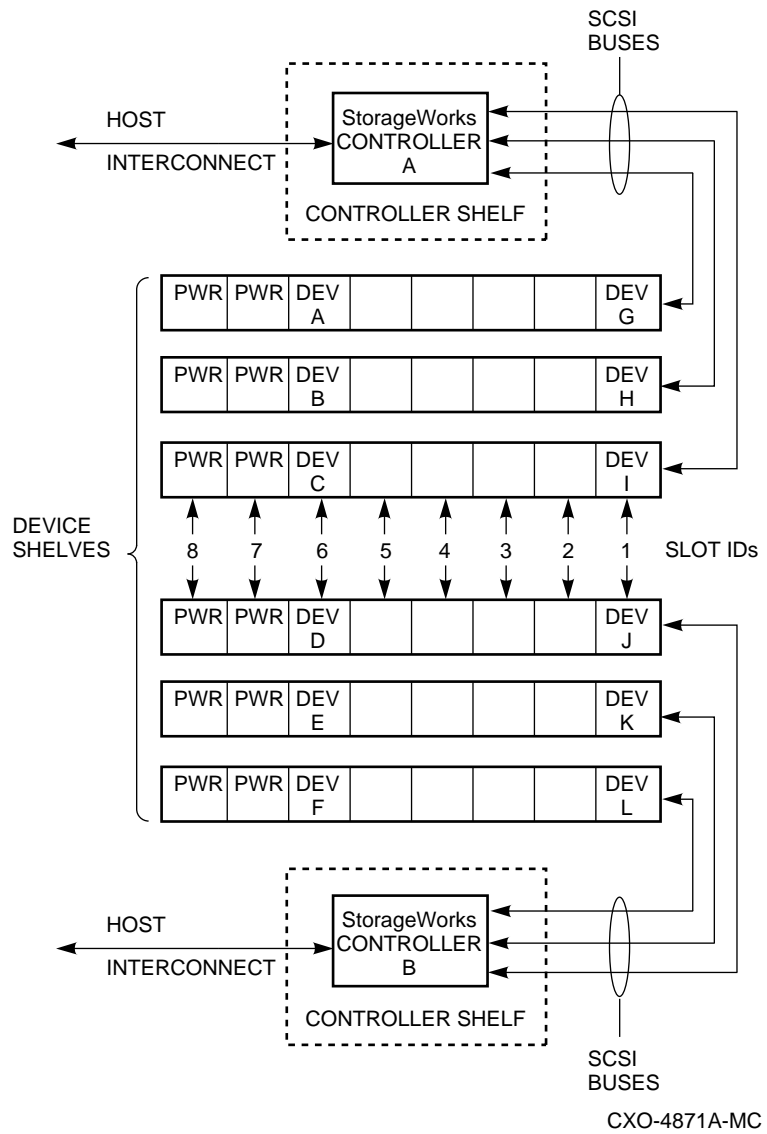
The use of independent controllers and redundant controllers is explained in Chapter **Error! Reference source not found.**

**Figure 1–1 Typical StorageWorks Configuration Using Redundant Controllers**



CXO-4870A-MC

**Figure 1–2 Typical StorageWorks Configuration Using Independent Controllers**



The final phase of the configuration process is implementing the configuration. Implementing the configuration occurs in the following stages:

- Ordering the components
- Preparing the site
- Installing the components
- Integrating the storage system
- Acceptance testing
- Populating the storage system with data

---

**Note**

---

This guide provides only a brief overview of migration planning. The overview appears in Appendix **Error! Not a valid filename.**

---

Table 1–1 describes the organization for the remainder of the guide.

**Table 1–1 Organization of the Configuration Guide**

<b>If you need information on ...</b>	<b>See ...</b>
Analyzing the storage system requirements	Chapter <b>Error! Not a valid filename.</b>
Planning the configuration	Chapter <b>Error! Not a valid filename.</b>
Selecting StorageWorks products	Chapter <b>Error! Not a valid filename.</b>
Implementing the StorageWorks configuration	Chapter <b>Error! Not a valid filename.</b>
General configuration rules	Appendix <b>Error! Not a valid filename.</b>
RAID overview	Appendix <b>Error! Not a valid filename.</b>
Migration planning overview	Appendix <b>Error! Not a valid filename.</b>

---

## Analyzing the Storage Requirements

Analyzing the storage requirements is the first phase of the configuration process. The goal of analysis is to understand the storage requirements, allowing you to later make cost effective purchasing decisions for StorageWorks products while providing solutions to the storage needs of users and applications that access the storage system.

After completing the analysis phase, you will have developed a list of requirements that defines the storage system based on needs defined in terms of the following:

- Users that will access the storage system
- Applications that will access the storage system
- Storage capacity
- Data categories of the stored data
- Applications' I/O requirements
- Data availability requirements

You may supplement the previous list with other criterion more specific to your needs. Combining your specific criterion with those listed will provide valuable information necessary to begin the planning phase discussed in Chapter 3.

Examples 2-1 and 2-2 show two worksheets, a User Analysis Worksheet and an Applications Analysis Worksheet. The worksheets are used throughout this chapter. The worksheets require that user groups and applications be entered separately to help identify similarities in the storage requirements of different user groups or of different applications. Similar storage requirements may permit user groups or applications to share a common storage configuration. Dissimilar storage requirements may indicate that multiple storage configurations are required as part of the StorageWorks solution.

---

### Note

---

Although the storage requirements for applications or user groups may be identical or similar enough to warrant the use of one storage configuration, security or other organizational requirements may preclude storing one organization's data in the same storage configuration as another organization's data.

---

## Case Study Overview

A fictional enterprise is used to provide examples for the configuration. The enterprise develops and markets software products and has the following functional organizations:

- Software development
- Software quality assurance
- Documentation development
- Office support

Other functional organizations would exist, but are not necessary for the examples used in this guide.

The software development organization creates new products. They keep hundreds of files in a central library system from which the files can be reserved for development work. When working on individual files, the developers reserve the files to their working directories. At the end of the business day each developer returns the files to the central library. Each evening the files are reserved by a batch job that merges the files, compiles the code, and tests the code for errors.

The software quality assurance organization uses the results of the previous evenings batch job to check for errors encountered in the executable part of the product.

The documentation development organization keeps another central library, separate from the software developers library, for documentation files. The documentation development organization shares the files to produce User Guides, Installation Guides, Release Notes, and other documentation that describes the software products. Like the software developers, the documentation developers reserve files from the central library to their working directories, update the files, then, when completed, return them to the central library.

The office support organization maintains contractual information on all customers and performs other office duties such as customer mailings and sales promotions.

Assume the user groups require the following storage:

- Software Development—5 megabytes (100 users times 500 kb per user)
- Quality assurance—1 megabyte (5 users times 200 kb per user)
- Document development—2 megabytes (10 users times 200 kb per user)
- Office support—2 megabytes (20 users times 100 kb per user)

Assume further, the enterprise runs applications having storage requirements as follows:

- Relational database—10 gigabytes
- Inventory control system—10 gigabytes
- Library management system—12 megabytes
- Computer aided design—15 gigabytes
- General file services—5 gigabytes

A final requirement is the total storage (approximately 40 gigabytes) must have a backup and restore policy.

These requirements are the basis for the configuration that is developed in this Guide.



**Example 2-1 User Analysis Worksheet**

<b>User Group—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>User Group—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>User Group—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>User Group—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	

**Example 2–2 Application Analysis Worksheet**

<b>Application—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Application—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Application—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Application—</b> _____				
Storage Capacity				
Active Data Category				
Dormant Data Category				
Inactive Data Category				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput				
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	

## **2.1 Analyzing Access to the Storage System**

Access to the storage system is by software applications. Some applications help users create, read, and modify data as well as perform general overhead tasks such as reading and responding to electronic mail messages. Other applications help users manipulate and transform data from one usable form to another. The number of users and applications must be determined in the analysis phase.

### **2.1.1 Determine the User Groups That Will Access the Storage System**

The users of the storage system can be grouped according to function. For example, the enterprise being analyzed in this discussion has software developers, a quality assurance team, office support, and documentation developers. Each user group has storage requirements that may be different from the other user groups.

Using the User Analysis worksheet, list each user group and the number of users in each group.

### **2.1.2 Determine the Applications That Will Access the Storage System?**

Every application has different requirements for storage. Those requirements will be noted as the analysis process continues.

Using the Application Analysis worksheet, list all the applications that will access the storage system.

## **2.2 Determine the Required Storage Capacity**

Storage capacity is the amount of storage, measured in megabytes or gigabytes, required to store:

- User data
- Application data
- Any executable images or temporary storage areas required by the applications

Determining the required storage capacity will assist in identifying the types and numbers of storage devices needed to satisfy the I/O and the data availability requirements.

Using both analysis worksheets, list the storage capacity requirements for each users group and for each application.

## 2.3 Identify the Data Categories

Data categories are classifications of data based on the frequency of data access. Each data category maps to one or more types of storage device, each of which is well suited for storage of data belonging to that data category.

Table 2–1 describes each data category and the recommended storage.

**Table 2–1 Data Category and Storage Type**

<b>Data Category</b>	<b>Recommended Storage</b>
<p><b>Active</b> Data that is frequently accessed and needs the fastest response time. Active data requires virtually immediate access, typically in milliseconds or less.</p>	<p><b>On-line</b> On-line storage does not require robotic or human intervention to make its data accessible. On-line storage is usually provided by solid state disks, magnetic disks, or other media that is logically and physically connected to the host system at all times.</p>
<p><b>Dormant</b> Data that is accessed less frequently and for which response time is less important. Dormant data can also be data that you have advance knowledge of when it will be required and can, therefore, plan when to make it active.</p>	<p><b>Near-line</b> Near-line storage should not require human intervention to make its data accessible. Near-line storage is usually provided by tape libraries, optical libraries or other robotic devices.</p>
<p><b>Inactive</b> Data that is not expected to be accessed frequently, but must be kept for business or legal purposes.</p>	<p><b>Off-line</b> Off-line storage can require human intervention to make its data accessible. Off-line storage is usually provided by Write-Once-Read-Many (WORM) or Read Only (RO) optical media that can be retrieved and made available in reasonable time frames.</p>

Initially, most data approaches being 100% active data. However, as the data matures, some of it migrates from the active category to the less demanding categories of dormant or inactive data. Over time the percentage of data in each category tends to stabilize. The 80/20 general rule states that on most systems, 80 percent of the I/O requests access only 20 percent of the stored data. The remaining 80 percent of the data is used infrequently. That remaining 80 percent is usually referred to as dormant or inactive data.

Digital advises that you plan for 100% active data during the configuration. However, expect to move the data between the three categories. As the data moves from the active category to the dormant or inactive category, expensive on-line storage can be made available to new users or applications by moving the dormant and inactive data to less expensive near-line or off-line storage.

Although the percentage of data in each category changes, it is worth while to estimate the percentage of an application's data that, over time, fits into each data category. Using the User Analysis and the Application Analysis Worksheets, list the percentage of each user group's and each application's data that belongs to the active, dormant, or inactive data category. Keep in mind the percentages will change over time and you will need to plan for the movement of data between data categories.

## 2.4 Determine Each Application's I/O Requirements

An application's I/O requirements are the I/O demands made of the storage system. The I/O requirements map to performance characteristics of different storage configurations. An application's I/O requirements, as considered in this guide, are divided into the following categories:

- I/O Activity
- Data throughput
- I/O response time

### 2.4.1 I/O Activity

I/O activity is the type of I/O performed by an application. I/O activity can be classified as any combination of the following:

- Data transfer-intensive
- I/O request-intensive
- Read-intensive
- Write-intensive

Data transfer-intensive applications must read or write large amounts of data per unit time, usually in a single stream. Applications such as computer assisted design, image processing, data collection and reduction, and sequential file processing tend to be data transfer-intensive.

I/O request-intensive applications execute large numbers of I/O requests per unit time. Applications such as interactive transaction processing, office automation, and general I/O or file services for timesharing tend to be I/O request-intensive.

Read-intensive applications perform a much higher percentage of read operations than write operations. Conversely, write-intensive operations perform a much higher percentage of write operations than read operations.

Using the User Analysis and the Applications Analysis Worksheets, list the I/O activity for each user group and each application that will access the storage system. Enter N/A (not applicable) if a user group or an application does not exhibit a tendency toward having any predominant transfer requirement.

### 2.4.2 Data Throughput

Data throughput is the sustainable transfer rate in bytes/sec between the storage media and the host system. Data throughput is primarily a concern for data transfer-intensive applications. If the storage system will support data transfer-intensive applications, contact the applications developers and have them supply the data throughput requirements for the applications.

Using the User Analysis and the Applications Analysis Worksheets, list the data throughput requirements for each user group and each application having data throughput requirements. Enter N/A if data throughput is not a primary requirement.

### 2.4.3 I/O Response Time

I/O response time is the interval between beginning and completing a data transfer request. For the purpose of this guide, an I/O response time of less than 50 ms is considered low, greater than 50 ms is considered high, and 50 ms is considered medium. I/O response time is primarily a concern for I/O request-intensive applications and is most affected by device access time.

Device access time is the time required for a storage device to move the requested data into position before a data transfer can begin. For example, a disk drive cannot begin to transfer data until it moves the read/write heads to the proper area of the disk and the disk rotates into the proper position.

Using the User Analysis and the Applications Analysis Worksheets, list the I/O response time requirements for each user group and each application having I/O response time requirements. Enter N/A if I/O response time is not a primary requirement.

## 2.5 Determine the Required Level of Data Availability

Data availability is the storage system's ability to provide a copy of the requested data on demand. Data availability is categorized as normal or extended in this guide.

Applications that require normal data availability are those that can tolerate some loss of service while one of the following occurs:

- The application's active data is restored from a backup copy
- The application's data is moved from a dormant or inactive state to the active state.

Applications that require extended availability are those that cannot tolerate a loss of service. When extended availability is required, you must plan a storage solution that includes device redundancy and/or a RAID solution.

Using the User Analysis and the Application Analysis Worksheets, list the data availability requirements (normal or extended) for each user group and application.

## 2.6 Summarizing the Analysis

The summary analysis combines, where possible, storage requirements. The summary analysis is based on the assumption that combining storage requirements allows for planning a more cost effective StorageWorks solution.

---

**Note**

---

Only those applications that have like requirements for I/O activity, data throughput, I/O response time, and data availability and have no conflicts with one another should have their storage requirements combined as discussed in this section. For example, if the files required by two or more applications with the same I/O requirements are placed on devices that meet the individual requirements and the applications are run simultaneously, the storage system may not perform as planned. Multiple applications running simultaneously and requesting the same storage can create I/O contention problems unless there has been proper planning for the increased I/O demands.

---

Using the Summary Worksheet, shown in Example 2–6, and the procedures described in Table 2–2 combine the data from the User Analysis and the Application Analysis Worksheets. Example 2–5 shows the results obtained after completing the Summary Worksheet using the Analysis Worksheets shown in Examples 2–3 and 2–4. Examples 2–3 and 2–4 are the Analysis Worksheets that would have resulted from an analysis of the fictional enterprise described at the beginning of this chapter.

**Table 2–2 Procedures for Completing the Summary Worksheet**

To summarize the...	perform the following steps:	For example, using the Worksheets in Examples 2–3 and 2–4, the ...
inactive data requirements	<ol style="list-style-type: none"> <li>For each entry on the Analysis Worksheets, multiply the values contained in the Storage Capacity row by the percentages in the Inactive Data Category row.</li> <li>Sum the values calculated in Step 1</li> <li>Enter the total on the Summary Worksheet in the Storage Capacity row for the Inactive Data Category.</li> </ol>	Relational database is $10\text{Gb} \times .06 = 0.6\text{ Gb}$ Inventory control is $10\text{ Gb} \times .05 = 0.5\text{ Gb}$ File services is $5\text{ Gb} \times .05 = 0.25\text{ Gb}$ $0.6\text{ Gb} + 0.5\text{ Gb} + 0.25\text{ Gb} = 1.35\text{ Gb}$ See Example 2–5.
dormant data requirements	<ol style="list-style-type: none"> <li>For each entry on the Analysis Worksheets, multiply the values contained in the Storage Capacity row by the percentages in the Dormant Data Category row.</li> <li>Sum the values calculated in Step 1</li> <li>Enter the total on the Summary Worksheet under the Storage Capacity column for the Dormant Data Category.</li> </ol>	Relational database is $10\text{ Gb} \times .75 = 7.5\text{ Gb}$ Inventory control is $10\text{ Gb} \times .70 = 7.0\text{ Gb}$ CAD is $15\text{ Gb} \times .80 = 12\text{ Gb}$ File services is $5\text{ Gb} \times .75 = 3.75\text{ Gb}$ $7.5\text{ Gb} + 7.0\text{ Gb} + 12\text{ Gb} + 3.75\text{ Gb} = 30.25\text{ Gb}$ See Example 2–5
active data requirements (initial)	<ol style="list-style-type: none"> <li>Beginning with the Summary Worksheet table row labeled Active and using as many rows as required, enter each unique combination of requirements that were defined during the analysis. The requirements to be combined are:                             <ul style="list-style-type: none"> <li>– I/O activity</li> <li>– Data throughput</li> <li>– I/O response time</li> <li>– Data availability</li> </ul> </li> <li>Using both Analysis Worksheets add the values contained in the Storage Capacity row for each user group and each application that matches a set of requirements entered in Step 1. (Note: This plans for 100 % active data as advised in Section 2.3.)</li> <li>In the Data Category row note which applications and user groups are being combined. Enter the results obtained in Step 2 under the appropriate Data Category of the Summary Worksheet in the Storage Capacity row.</li> </ol>	Examples 2–3 and 2–4 show the following unique combinations. <ul style="list-style-type: none"> <li>• All user groups have identical I/O and availability requirements</li> <li>• The relational database and inventory control applications have identical I/O and availability requirements.</li> <li>• The CAD application’s I/O and availability requirements are unique to the application.</li> <li>• The file services application’s I/O and availability requirements are unique to the application.</li> </ul> User groups total $5\text{ Mb} + 1\text{ Mb} + 2\text{ Mb} + 2\text{ Mb} = 10\text{ Mb}$ Customer database and inventory control $10\text{ Gb} + 10\text{ Gb} = 20\text{ Gb}$ CAD $15\text{ Gb}$ File services $5\text{ Gb}$ See Example 2–5

Continued on next page



**Table 2-2 (Cont.) Procedures for Completing the Summary Worksheet**

<b>To summarize the...</b>	<b>perform the following steps:</b>	<b>For example, using the Worksheets in Examples 2–3 and 2–4, the ...</b>
<p>active data requirements (projected)</p>	<ol style="list-style-type: none"> <li>1. Beginning with the Summary Worksheet table row labeled Active and using as many rows as required, enter each unique combination of requirements that were defined during the analysis. The requirements to be combined are: <ul style="list-style-type: none"> <li>– I/O activity</li> <li>– Data throughput</li> <li>– I/O response time</li> <li>– Data availability</li> </ul> </li>   <li>2. For each entry on the Analysis Worksheets, multiply the values contained in the Storage Capacity row by the percentages in the Active Data Category row.</li>   <li>3. Add the active storage requirements for each application and/or user group that matches a set of requirements entered in Step 1.</li>   <li>4. Enter the results (use parenthesis) on the appropriate Storage Capacity row of the Summary Worksheet.</li> </ol>	<p>Examples 2–3 and 2–4 show the following unique combinations.</p> <ul style="list-style-type: none"> <li>• All user groups have identical I/O and availability requirements</li> <li>• The relational database and inventory control applications have identical I/O and availability requirements.</li> <li>• The CAD application’s I/O and availability requirements are unique to the application.</li> <li>• The file services application’s I/O and availability requirements are unique to the application.</li> </ul> <p>User groups  Software Dev. 5 Mb X 1.00 = 5 Mb  Quality Assurance 1 Mb X 1.00 = 1 Mb  Document Dev. 2 Mb X 1.00 = 2 Mb  Office Support 2 Mb X 1.00 = 2 Mb</p> <p>Applications  Relational database  10 Gb X .19 = 1.9 Gb  Inventory control  10 Gb X .25 = 2.5 Gb  CAD  15 Gb X .20 = 3.0 Gb  File services  5 Gb X .20 = 1.0 Gb</p> <p>User groups total  5 Mb + 1 Mb + 2 Mb + 2 Mb = 10 Mb</p> <p>Customer database and inventory control  1.9 Gb + 2.5 Gb = 4.4 Gb</p> <p>CAD  3.0 Gb</p> <p>File services  1.0 Gb</p> <p>See Example 2–5</p>

The information in the Summary Worksheet is used in Chapters 3 and 4 to plan the configuration and to select the StorageWorks products.

**Example 2–3 Completed User Analysis Worksheet for the Fictional Company**

<b>User Group—Software Development (100 users)</b>				
Storage Capacity	5000 KB			
Active Data Category	100			
Dormant Data Category	0			
Inactive Data Category	0			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	<b>N/A</b>
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
<b>User Group—Quality Assurance (5 users)</b>				
Storage Capacity	1000 KB			
Active Data Category	100			
Dormant Data Category	0			
Inactive Data Category	0			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	<b>N/A</b>
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
<b>User Group—Document Development (10 users)</b>				
Storage Capacity	2000 KB			
Active Data Category	100			
Dormant Data Category	0			
Inactive Data Category	0			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	<b>N/A</b>
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
<b>User Group—Office Support (20 users)</b>				
Storage Capacity	2000 KB			
Active Data Category	100			
Dormant Data Category	0			
Inactive Data Category	0			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	<b>N/A</b>
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	

**Example 2–4 Completed Application Analysis Worksheet for the Fictional Company**

<b>Application—Relational Database</b>				
Storage Capacity	10 GB			
Active Data Category	19			
Dormant Data Category	75			
Inactive Data Category	6			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
<b>Application—Inventory Control</b>				
Storage Capacity	10 GB			
Active Data Category	25			
Dormant Data Category	70			
Inactive Data Category	5			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
<b>Application—CAD</b>				
Storage Capacity	15 GB			
Active Data Category	20			
Dormant Data Category	80			
Inactive Data Category	0			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
<b>Application—File Services</b>				
Storage Capacity	5 GB			
Active Data Category	20			
Dormant Data Category	75			
Inactive Data Category	5			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	

**Example 2–5 Completed Summary Worksheet for the Fictional Company**

<b>Data Category—Inactive (Near Line)</b>				
Storage Capacity	1.35 GB			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—Dormant (Off-Line)</b>				
Storage Capacity	30.25 GB			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—Active (All Users)</b>				
Storage Capacity	10 Mb (10 Mb)			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
<b>Data Category—File Services</b>				
Storage Capacity	5.0 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	

**Example 2–6 Summary Worksheet**

<b>Data Category—Inactive (Near Line)</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—Dormant (Off-Line)</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—Active (All Users)</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
<b>Data Category—</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	



---

## Planning the StorageWorks Configuration

This chapter helps you plan the StorageWorks configuration based on the analysis completed in Chapter 2. Planning the configuration requires an examination of the requirements defined in the analysis phase, then choosing a solution that meets those requirements.

Planning is divided into three major areas:

- Planning for storage of active data
- Planning for storage of dormant and inactive data
- Planning for backup and restore operations

Example 3–1 shows the Configuration Planning Worksheet. The worksheet assists in selecting StorageWorks products that satisfy the requirements defined in the analysis phase. (Note, the Configuration Planning Worksheet is an extension of the Summary Worksheet prepared in Chapter 2.) Table 3–1 describes how the Configuration Planning Worksheet information is used to assist in the selection of StorageWorks products.

**Table 3–1 Configuration Planning Worksheet and Product Selection**

The requirements for ...	will assist in selecting the...
data category	type of storage (on-line, near-line, or off-line).
I/O activity	number of controllers.
data throughput	type of devices, the type of shelves, the controller and device configurations.
I/O response time	storage devices.
data availability	degree of device redundancy and/or RAID solution.

---

### Note

---

Examples of worksheet use show only the parts of interest for the topic under discussion. The completed worksheet for the planning phase is shown in Example 3–5

---

**Example 3–1 Configuration Planning Worksheet**

<b>Data Category—Inactive (Near Line)</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Dormant (Off-Line)</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Active (All Users)</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				



**Example 3-1 (Cont'd) Configuration Planning Worksheet**

<b>Data Category—</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended		N/A
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—</b>				
Storage Capacity				
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended		N/A
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				

## 3.1 Planning for Storage of Active Data

Active data requires planning for the following:

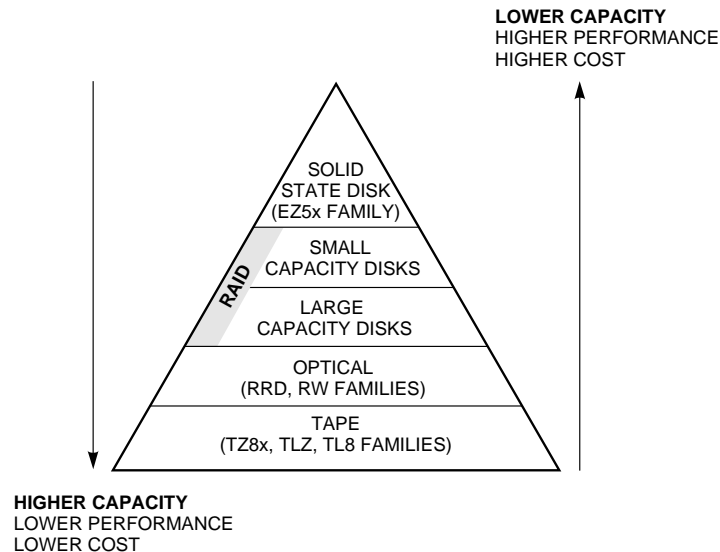
- Capacity
- I/O performance
- Data availability
- RAID solutions to improve performance and availability

### 3.1.1 Planning for Capacity

Storage capacity is the one requirement you must satisfy. Planning for capacity helps you to cost-effectively choose the devices best able to meet your storage capacity requirements. You can satisfy the storage capacity requirement by distributing it across many devices. Which devices you choose should be influenced by device type, cost/Mb, and performance.

The relationship among these factors is illustrated in Figure 3–1. In general, high-performance devices have a lower capacity and higher cost/Mb than lower performance high-capacity devices. High-capacity devices trade performance for the ability to store large amounts of data.

**Figure 3–1 Storage Relationships**



Digital advises that when determining your storage requirements you consider your future storage requirements as well. The storage capacity required by most applications tends to grow over time. Similarly, the number of users tends to increase over time. As the number of users increases, the storage requirements increase.

Early planning for future storage capacity requirements will save time and cost. For example, you may choose a larger enclosure than required, but not completely populate it with storage devices. Or, you may decide to use multiple enclosures when one would suffice, partially populating each enclosure with storage devices. In either situation, the power, cooling, and service area requirements will, probably, change. Planning for future capacity helps to ensure easy and efficient additions to the storage system.

## 3.1.2 Planning for Storage System Performance

Performance is a relative measure of the extent to which the I/O requirements of data throughput, I/O response time, and type of I/O meet the requirements defined during the analysis phase. Planning to meet the I/O requirements of the storage system requires knowledge of the various configuration options that affect performance. Once you decide on the configuration that best meets your requirements, then you can select the component parts that, together, will meet those performance requirements

When planning for storage system performance consider:

- Controller caching
- Multiple independent controllers
- 16-bit SCSI shelves

---

### Note

---

For high performance systems, you should also consider RAID configurations. See Section 3.1.4.

---

### 3.1.2.1 Controller Caching

Controller caching improves performance by using a high speed memory to buffer the data transfers between the controller and the host system. The buffering allows data transfers with the host to occur at the speed of the cache memory system and not at the slower device transfer speeds. Most StorageWorks controllers support the following types of cache:

- Read
- Write-back

Read cache improves performance during read operations by keeping recently accessed data blocks and adjacent data blocks in cache memory. Subsequent requests for data still in cache do not require a device access. The request completes at the maximum bus speed using the copy contained in the cache memory.

Write-Back cache improves performance during write operations by allowing host processors to continue working rather than waiting for the write operations to be completed to the storage device. Using a write-back cache can provide a response time reduction of 50% or more. Write-back cache boosts the throughput of all applications that write to storage devices. When write-back cache is activated, read cache is automatically brought on-line.

When caching options are used and the controller simultaneously caches both read and write requests, controllers, such as the DSSI-to-SCSI, that dynamically allocate cache based on real-time read/write activity, provide significant performance gains over controllers that do not provide caching.

If you plan to use controller caching, use the Controller Configuration rows of the Configuration Planning worksheet to note the type of caching desired. For example, if both types of caching were desired with the controller that will provide access to the storage required for the relational database, inventory control, and CAD applications of the fictional enterprise, the worksheet would be filled out as shown in Example 3-2 (See lightly shaded area). In rows with multiple choices listed, the choice for our example is shown in bold type.

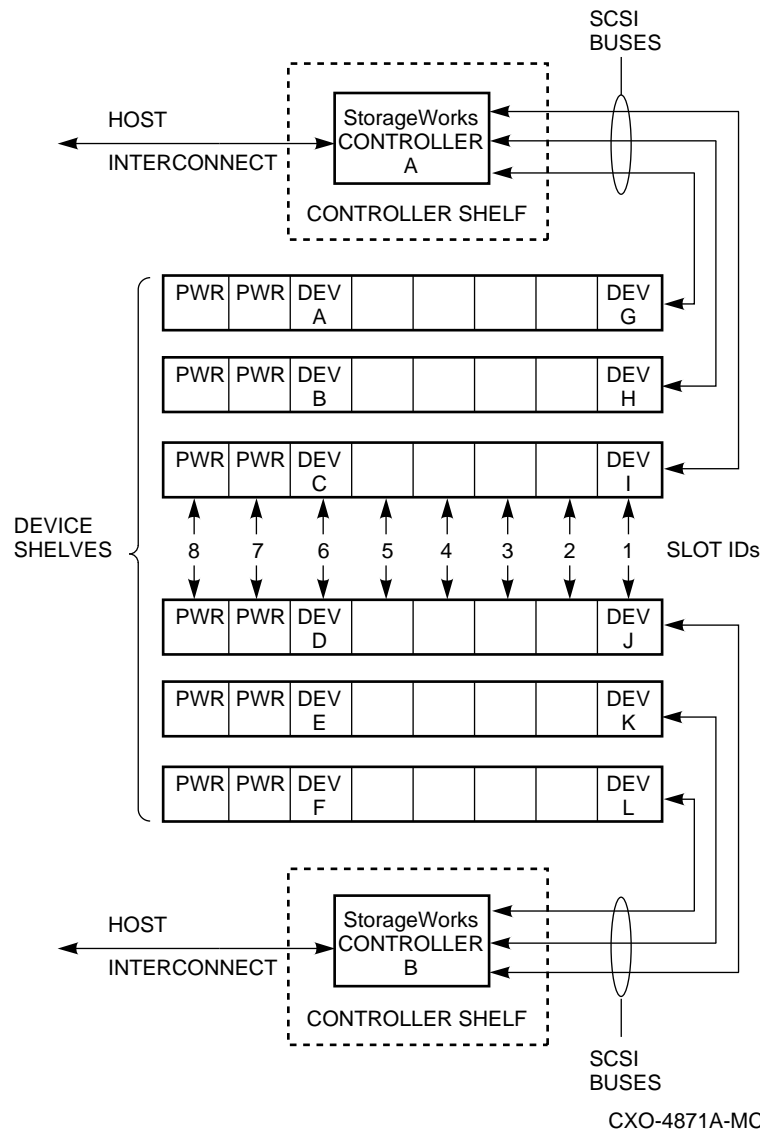
**Example 3–2 Configuration Planning Worksheet—After Adding Performance Information**

<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache			
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache		Independent controllers	
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration	Independent controller (use one of the CAD controllers)			
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				

### 3.1.2.2 Multiple Independent Controllers

Multiple independent controllers, as discussed in this guide, are two or more controllers that work together to satisfy the I/O requirements of one or more common applications, but do not share access to the same data. Each controller provides access to only a part of the application's data. Figure 3-2 shows a simple configuration that distributes an application's data across six disks (Devices A through F) using two 3-port StorageWorks controllers to provide access to one disk per port. StorageWorks controller A can access only the data on devices A, B and C. StorageWorks controller B can access only the data on devices D, E, and F.

**Figure 3-2 Independent Multiple Controllers Configuration**



Although the purpose of using multiple independent controllers is to distribute the workload of one or more applications, each controller, because it is independent of the other, can service the workload requests of other applications that do not require the workload to be distributed. Figure 3-2 shows that StorageWorks controller A provides an access path to

devices G, H, and I, and StorageWorks controller B provides an access path to devices J, K, and L. Each group of devices could contain the data of separate applications.

If you plan to use multiple independent controllers, use the Controller Configuration rows of the Configuration Planning worksheet to note your decision. For example, assume the fictional enterprise will use two independent controllers to provide access to the storage of the CAD application. Further assume, that one of the controllers will provide access to the storage of the file services application. The worksheet would be filled out as shown in Example 3–2 (See heavily shaded area). In rows with multiple choices listed, the choice for our example is shown in bold type.

### **3.1.2.3 Using 8 or 16-Bit SBB Shelves**

SBB shelves are available with an 8-bit or a 16-bit SCSI bus. The use of a shelf containing a 16-bit SCSI bus supports twice the data transfer rate as an SBB shelf containing an 8-bit SCSI bus. To take advantage of the higher transfer rates, you must use storage devices and controllers that are capable of transferring 16-bits at a time. Not all StorageWorks controllers and devices support 16-bit transfers. This does not prevent you from using the 16-bit shelves with other StorageWorks devices and controllers. The 16-bit shelves can be configured to perform 8-bit or 16-bit transfers. For future growth Digital recommends that you use 16-bit SBB shelves whenever possible.

### 3.1.3 Planning for Data Availability

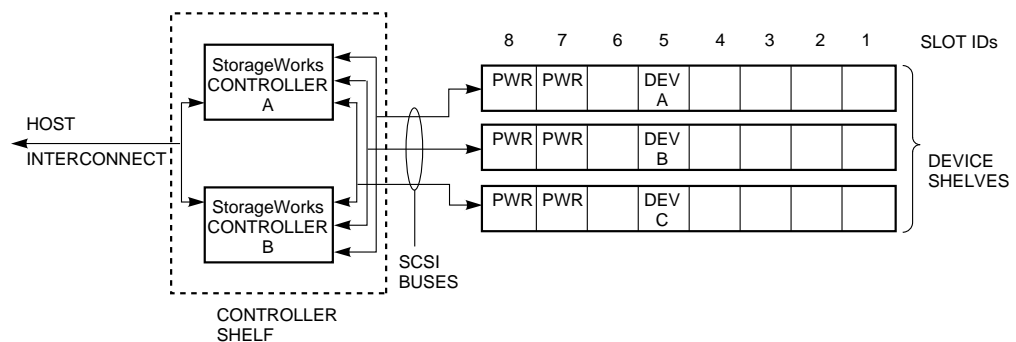
Planning for data availability ensures the data is available when required. If your storage system requirements can tolerate periods when the data is unavailable, then the use of good storage backup policies and fast restores of data are probably adequate. If, on the other hand, your data availability requirements are extended, you can use one or more the following solutions:

- Using redundant controllers
- Using redundant power supplies
- Using a RAID solution (See Section 3.1.4)

#### 3.1.3.1 Using Redundant Controllers

Redundant controllers improve data availability because each controller provides access to the same physical storage. A side benefit is the increased performance realized when the controllers share the I/O load associated with the storage. When one controller fails, the remaining controller assumes the I/O workload of the failed controller as well as its own workload. Although performance is reduced until the failed controller is replaced, the data remains available. Figure 3–3 shows a configuration using redundant controllers.

**Figure 3–3 Redundant Controller Configuration**



CXO-4870A-MC

If you plan to use redundant controllers, use the Controller Configuration rows of the Configuration Planning worksheet to note your decision. For example, if redundant controllers were desired to provide access to the storage required for the relational database and the inventory control applications of the fictional enterprise, the worksheet would be filled out as shown in Example 3–3 (See lightly shaded area). In rows with multiple choices listed, the choice for our example is shown in bold type.

#### 3.1.3.2 Using Redundant Power Supplies

Redundant power supplies improve data availability by providing continuous power to controllers and storage devices if one power supply fails. Redundant power supplies are available for device, controller, and combination shelves. In larger enclosures, such as the SW800 Data Center cabinet, you can also provide redundant Cable Distribution Units to provide an additional level of power system redundancy.

If you plan to use redundant power supplies, use the Power Supply rows of the Configuration Planning worksheet to note your decision. For example, if redundant power supplies were desired for all on line storage of the fictional enterprise, the worksheet would be filled out as shown in Example 3–3 (See heavily shaded area).

**Example 3–3 Configuration Planning Worksheet—After Adding Data Availability Information**

<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache		Redundant controllers	
Power Supply	Redundant power supplies			
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration				
Power Supply	Redundant power supplies			
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration				
Power Supply	Redundant power supplies			
RAID Arrays				
Storage Devices				
Shelves/Enclosures				



### 3.1.4 Using RAID to Improve Performance and Availability

Review Appendix B if you are not familiar with RAID technology. It provides a brief discussion of the RAID levels supported by StorageWorks products.

---

**Note**

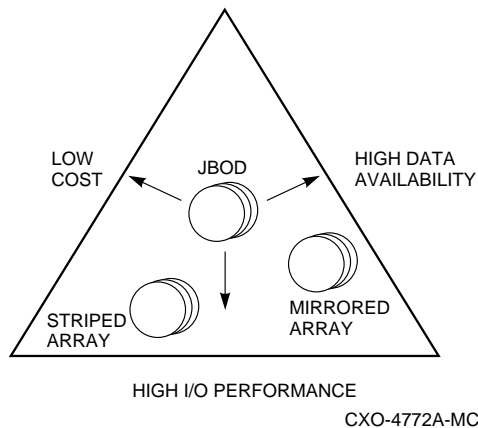
---

Configuring a storage system using RAID for availability does not replace a defined and enforced backup policy. Always have and enforce a backup policy.

---

RAID improves performance and availability but at increased cost. Figure 3–4 shows that a striped array has higher performance and approximately the same cost as multiple independent disks (sometimes call JBOD or “Just a Bunch of Disks”). The higher performance is achieved at the expense of lower data availability because if any array member fails, data for the entire array is lost and must be restored from backup media. Similarly the mirrored array has higher performance and data availability than multiple independent disks, but at an increased cost. Comparing the striped array to the mirrored array shows the striped array has higher performance and lower cost than the mirrored array, but the mirrored array provides higher data availability. You can improve any two factors, but always at the expense of the remaining factor.

**Figure 3–4 Cost–Availability–Performance Triangle**



If you are planning to use RAID in your configuration, use the RAID Arrays rows of the Configuration Planning worksheet to note the RAID level or, if RAID is not a part of the solution, enter none. For example, if Parity RAID (RAID Level 5) were desired for the relational database and inventory control applications, Mirroring (RAID Level 1) were desired for the CAD application, and no RAID solution for the file services application of the fictional enterprise, the worksheet would be filled out as shown in Example 3–4 (See lightly shaded area). In rows with multiple choices listed, the choice for our example is shown in bold type.

#### Using Sparesets

All RAID configurations, except Striping (RAID Level 0), provide protection against loss of data due to a single device failure. However, there is no protection against the failure of a

second device. If a second device fails before the first is repaired or replaced, the data will be lost and remain unavailable until the data can be restored from back-up media. To protect against this possibility, plan a spareset for the RAID array.

A spareset is a pool of spare drives that can be used as both an automatic replacement and a manual replacement for RAID array members that have been removed. Because sparesets can be associated with many RAID arrays, you need only one or two drives in a spareset to ensure that the RAID arrays continue to operate with redundancy if a member fails. The minimum spareset is one drive, and if that spareset is shared by multiple RAIDsets, the spareset drive must have the same capacity as the largest drive used in any RAIDset.

---

**Note**

---

A spareset cannot be used in a RAID Level 0 configuration since there is no inherent redundancy from which to restore the data if a disk fails.

---

If you plan to configure a spareset, use the RAID Arrays rows of the Configuration Planning worksheet to note the requirement for a spareset. For example, if a spareset were desired for the relational database and inventory control applications of the fictional enterprise, the worksheet would be filled out as shown in Example 3-4 (See heavily shaded area).

**Example 3–4 Configuration Planning Worksheet—After Adding RAID Information**

<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache		Redundant controllers	
Power Supply	Redundant power supplies			
RAID Arrays	RAID Level 5		Sparesets	
Storage Devices				
Shelves/Enclosures				
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration				
Power Supply	Redundant power supplies			
RAID Arrays	RAID Level 1			
Storage Devices				
Shelves/Enclosures				
<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration				
Power Supply	Redundant power supplies			
RAID Arrays	None			
Storage Devices				
Shelves/Enclosures				

### 3.1.5 Tradeoffs Between Performance and Availability

Obtaining the highest performance while, at the same time, maintaining the highest availability in the same configuration is not a realistic goal if cost is paramount. Due to cost constraints, you usually give up some performance requirements to achieve higher availability or vice-versa.

Some of the tradeoffs are as follows:

- Using individual disks (“Just a Bunch of Disks”) will provide less performance and less availability than RAID in most situations.
- Using a RAID Level 5 and a spare member will cause a temporary decrease in performance after the loss of a member disk.
- Using RAID Level 0 will provide higher performance than an equivalent number of non-arrayed disks. However, if any member of the array fails the data stored across the entire array is lost. The data remains unavailable until it has been restored from backup media.

## 3.2 Planning for Storage of Dormant and Inactive Data

Dormant and inactive data should be stored on media that is near-line or off-line. Both data types have different storage requirements than active data that is stored using on-line media. In both cases access to the data usually requires human or robotic intervention to move the media to an on-line state. The primary concern is the type of devices used to store the media and a policy for refreshing, if required, the data at defined intervals.

---

#### Note

---

Typically, the I/O requirements for data that is brought on-line from near-line or off-line media are low and can be satisfied by the storage device used to make the near-line or off-line media available to the host system. If the data has higher I/O requirements than can be handled by the storage device, the data must be copied to another on-line device such as a magnetic disk.

---

Using the Configuration Planning Worksheet, note the storage requirements for inactive and for dormant data. During the device selection phase (discussed in Chapter 4) selecting one or more devices is necessary to accommodate the storage requirements.

## 3.3 Planning for Backup and Restore Operations

Every storage system must have a backup and restore strategy to safeguard against permanent data loss. To implement the strategy one or more devices that can store the amount of data to be backed up should be planned as part of the storage configuration.

Although back-up data is typically kept off line, the data requirements for back-up data and dormant data are different enough to justify different storage devices for each. Storage devices best suited for backing up and restoring data are tape loaders that operate unattended.

Using the Configuration Planning Worksheet determine the amount of data to be backed up by adding the storage capacity requirements for all the data categories. During the device selection phase (discussed in Chapter 4) selecting one or more devices is necessary to accommodate the backup requirements.

**Example 3–5 Configuration Planning Worksheet—After Completing the Planning Phase**

<b>Data Category—Inactive (Near Line)</b>				
Storage Capacity	1.3GB			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Dormant (Off-Line)</b>				
Storage Capacity	30.25			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Active (All Users)</b>				
Storage Capacity	10 MB (10 MB)			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache; redundant controllers			
Power Supply	Redundant power supplies			
RAID Arrays	RAID Level 5; Spareset			
Storage Devices				
Shelves/Enclosures				

**Example 3-5 (Cont'd) Configuration Planning Worksheet—After Completing the Planning Phase**

<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache; Independent controllers			
Power Supply	Redundant power supplies			
RAID Arrays	RAID Level 1			
Storage Devices				
Shelves/Enclosures				
<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration	Independent controllers			
Power Supply	Redundant power supplies			
RAID Arrays	None			
Storage Devices				
Shelves/Enclosures				

---

## Selecting StorageWorks Products

This chapter provides basic guidelines for selecting StorageWorks products based on the analysis and planning accomplished in the previous chapters

Once you have analyzed your storage system requirements and planned the configuration you can begin selecting from the following StorageWorks products to complete the configuration:

- Controllers and adapters
- Storage Devices
- Shelves
- Power components
- Enclosures
- Cables and accessories

---

**Note**

---

The devices selected in the examples were valid at the time this Guide was published. But, please consult the *StorageWorks Products Catalog* or the *Systems and Options Catalog* for the latest information on StorageWorks products

---

The Configuration Planning Worksheet is used throughout this chapter to assist in the selection of the products. Examples of worksheet use show only the parts of interest for the topic under discussion. The completed worksheet for the selection phase is shown in Example 3–5.

## 4.1 Selecting Controllers and Adapters

Each StorageWorks controller provides an interface to one, and only one, host I/O bus or interconnect and supports 1, 3, or 6 SCSI device buses. StorageWorks controllers may be connected to foreign host I/O buses or interconnects by adapters.

Typical I/O buses or interconnects and the StorageWorks products that interface to them include the following:

- CI (Computer Interconnect)—HSJ family of controllers
- DSSI—HSD family of controllers
- SCSI (Small Computer System Interface)—HSZ family of controllers
- XMI (eXtended Memory Interconnect)—KFMSA adapter and HSZ family of controllers
- EISA—KZESA adapter and HSZ family of controllers
- QBUS—KFQSA adapter and HSZ family of controllers
- PCI—KZPSA adapter and HSZ family of controllers

Select the controller and, if required, the adapter based on the following:

- Support of the host I/O bus or interconnect
- Number of SCSI ports
- Supported RAID levels (if configuring a RAID solution)

Because controllers may not support every available storage device, note the storage devices with which the controller is compatible. You will need the information when selecting storage devices (described in Section 4.2).

Refer to the *StorageWorks Products Catalog* or the *Systems and Option Catalog* for information to help you select the StorageWorks controller and, if required, the adapter.

Enter the controller and, if required, the adapter types in the Controller Configuration rows of the Configuration Planning worksheet. Using the example enterprise, if HSJ40 controllers were selected to provide access to all the required storage, the worksheet would be filled out as shown in Example 4-1 (See lightly shaded area).



**Example 4–1 Configuration Planning Worksheet—After Adding Controller and Device Selections**

<b>Data Category—Active (All Users)</b>				
Storage Capacity	10 MB (10 MB)			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration	HSJ40			
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache; redundant controllers			HSJ40s
Power Supply				
RAID Arrays				
Storage Devices	RZ28–VA, 16 drives (4 RAIDsets); RZ28–VA, 4 drives (Spareset)			
Shelves/Enclosures				
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Independent controllers			HSJ40s
Power Supply				
RAID Arrays				
Storage Devices	RZ28–VA, 18 drives (3 RAIDsets)			
Shelves/Enclosures				

**Example 4–1 (Cont’d) Configuration Planning Worksheet—After Adding  
Controller and Device Selections**

<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration	Independent controllers (Use one of the CAD controllers)			HSJ40
Power Supply				
RAID Arrays				
Storage Devices	RZ28–VA, 3 drives			
Shelves/Enclosures				

## 4.2 Selecting Storage Devices

Selecting storage devices is a two step process: selecting the types of devices and selecting the number of each device.

### 4.2.1 Selecting the Device Types

Select storage device types based on the following:

- Data category (active, dormant, or inactive) of the data to be stored
- I/O requirements for data access
- Capacity of the device

Each type of storage device has characteristics that make it more suitable than others for storage of data based on the data category and on the I/O requirements for data access. StorageWorks devices provide the following types of storage devices:

- Solid state disk drives
- Magnetic disk drives (large and small capacity)
- Optical drives (Magneto-optic, Write Once Read Many, and CD-ROM drives)
- Tape drives, loaders, and libraries

Solid state disks are best suited for storage of the most active and frequently accessed data. Typically, the most active and frequently accessed data is 1 to 3% of the storage system's data.

Magnetic disk drives are best suited for storage of active data. Magnetic disk drives are available in data capacities ranging from 4.25 Megabytes to 4.3 Gigabytes per spindle with media transfer rates from less than 1 Megabyte per second to more than 7 megabytes per second.

Optical disk drives and CD-ROM drives provide low cost storage for bulk data.

Tape drives and tape libraries are best suited for inactive or dormant data.

Tape Loaders are best suited for unattended transfers of large amounts of data such as is associated with backup or restore operations.

In general, the following is true of StorageWorks storage devices:

- Smaller capacity disks provide better access times than larger capacity disks and have a higher cost/Mb.
- Larger capacity disks generally offer better data transfer rates and lower cost/Mb than smaller capacity disks
- Disks that rotate at a higher RPM perform better than disks that rotate at a lower RPM

For data transfer intensive applications, select devices with high data transfer rates. For I/O request intensive applications, select devices with a fast access time.

Refer to the *StorageWorks Products Catalog* or the *Systems and Options Catalog* to select the StorageWorks devices to be used in the configuration. Check for compatibility between the devices you select now and the controller you selected in Section 4.1.

For example, a TLZ6L would meet the requirements to backup the 40 Gbytes of storage in the fictional enterprise and can be mounted in a StorageWorks enclosure. The tape loader provides unattended operations and provides backup performance of 1 Gbyte per hour.

---

**Note**

---

Although the backup performance of 1 Gbyte per hour requires 40 hours to backup the full storage of the fictional enterprise, it is assumed that full backups are scheduled during long weekends. Incremental backups, which backup only the data that has changed, could easily be accomplished on a nightly basis.

---

## 4.2.2 Determining the Number of Storage Devices

The number of storage devices is determined by the total capacity required, the capacity of the individual devices, and for magnetic disks, whether or not a RAID solution is part of the device configuration.

For all configurations except those using RAID solutions, divide the total storage capacity required by the storage capacity of the selected device and round up, if required, to the next whole number. That will be the number of devices required. For example, using RZ28-VA disk drives having a capacity of 2.05 GB per drive would require three drives to store the 5 Gb of storage required by the file services applications of the fictional enterprise.

If a RAID configuration is part of the storage solution, one decision will be the number of member disks in an array. When deciding the number of member disks consider the following tradeoffs:

- For a given RAIDset capacity, RAIDsets using many smaller disks give better performance than RAIDsets using fewer larger disks. This is because with smaller disks, there are more members in the RAIDset to distribute the I/O workload.
- For a given workload capacity, fewer RAIDsets perform better than more RAIDsets. This is because minimizing the number of RAIDsets minimizes the workload imbalance among RAIDsets.
- The greater the number of RAIDset members, the longer the regeneration and reconstruction times when a RAID member disk fails.
- During regeneration and reconstruction, the larger the number of members in the RAIDset, the greater the performance impact on all applications running on that RAIDset.

---

**Note**

---

Although fewer members in a RAIDset increase data availability, the large cost of increasing the number of power supplies and controllers to accommodate the extra disks needed for smaller RAIDsets may be prohibitive.

---

For Mirroring (RAID Level 1) solutions, additional storage must be provided to account for the data redundancy. Perform the following steps to determine the required number of magnetic disks:

- Determine the number of RAIDsets across which the data will be stored.
- Determine the storage capacity of each RAIDset.

- For each RAIDset, calculate the number of disks required to provide the storage capacity. It is strongly recommended that all member disks of a RAIDset be of the same type and size. Divide the RAIDset capacity by the capacity of the smallest disk in the RAIDset. Round up, if required, to the next whole number
- For each RAIDset, multiply the number of disks by the level of data redundancy. For example, a single level of data redundancy requires multiplication by two; two levels of data redundancy require multiplication by three.

Using the example enterprise, assume that three RAIDsets each providing 5 GB of storage and using RAID Level 1 will be used to meet the requirements of the CAD application. Each RAIDset would require six RZ28-VA disks with a capacity of 2.05 GB per drive. The total number of drives required would be 18.

For Parity RAID (RAID Level 5) solutions, additional storage must be added to account for the parity data used to reconstruct data when a RAID member fails. When determining the number of magnetic disks, perform the following steps:

- Determine the number of RAIDsets across which the data will be stored.
- Determine the required storage capacity of each RAIDset.
- For each RAIDset, calculate the number of disks required to provide the storage capacity. Divide the RAIDset capacity by the capacity of the smallest drive in the RAIDset. Round up, if required, to the next whole number.
- For each RAIDset, add one additional disk to the number calculated above.

Using the example enterprise, assume that four RAIDsets, each providing 5 GB of storage and using RAID 5, will be used to meet the requirements of the relational database and inventory control applications. Each RAIDset would require four RZ28-VA disks with a capacity of 2.05 GB per drive. The total number of drives required would be 16.

The use of sparesets in a RAID solution must also be included when calculating the number of required storage devices. Using the example enterprise, if a four-drive spareset is used to provide an extra measure of data safety for the RAIDsets, then four additional drives are required. The drives must be of the same or larger capacity as the smallest capacity RAIDset member drive.

Enter the number of storage devices in the Storage Devices rows of the Configuration Planning worksheet. Using the examples from this section, the worksheet would be filled out as shown in Example 4-1 (See heavily shaded area).

### 4.3 Selecting Shelves

Three types of StorageWorks shelves are available: controller shelves, device shelves, and RAID shelves. Table 4–1 specifies the maximum capacities of each shelf.

**Table 4–1 StorageWorks Shelf Capacities**

Shelf Type	Maximum Capacity			
	3.5 inch devices	5.25-inch devices	Controllers*†	Power Supplies
8-bit device shelf	6 or 7	2	1*	2
16-bit device shelf	6 or 7	2	1*	2
7 device RAID shelf	6 or 7	2	2	2
24 device RAID shelf	24	8	2	8
Controller shelf	N/A	N/A	2	2
* Only the HSD05 or the HSD10 can be installed in device shelves † All controllers except the HSD10 and HSD15 can be installed in the RAID shelves or the controller shelf				

### 4.3.1 Device Shelves

Figure 4–1 shows a typical device shelf. Device shelves are normally used when one of the combination shelves (described in Section 4.3.3) is unsuitable or when making small additions to an existing configuration that has no available space for installing storage devices. If device shelves are selected, one or more controller shelves may also be needed for the configuration. The number of device shelves depends on the following:

- The number of storage devices
- The form factor of the storage devices

**Figure 4–1 Typical Device Shelf**




Table 4–2 shows the required number of device shelves for 3.5-inch devices connected to controllers having 3 or 6 SCSI ports. Table 4–3 shows the required number of device shelves for 5.25-inch devices connected to controllers with 3 or 6 SCSI ports. Calculate the total number of device shelves for each controller in your configuration.

**Table 4–2 Device Shelf Requirements for 3.5 Inch Devices**

If the number of SCSI Ports is...	and the number of devices is...	Then the minimum number of device shelves required is...
3	1 to 2	1
	3 to 12	2
	13 to 18	3
6	1 to 2	1
	3 to 4	2
	5 to 18	3
	19 to 24	4
	25 to 30	5
	31 to 36	6

**Table 4–3 Device Shelf Requirements for 5.25 Inch Devices**

<b>If the number of SCSI Ports is...</b>	<b>and the number of devices is...</b>	<b>Then the minimum number of device shelves required is...</b>
3 or 6	1 to 2	1
	3 to 4	2
	5 to 6	3
	7 to 8	4
	9 to 10	5
	11 to 12	6
6	13 to 14	7
	15 to 16	8
	17 to 18	9
	19 to 20	10
	21 to 22	11
	23 to 24	12

Using the example enterprise, if the storage for the CAD and the file services applications is to be accessed through two independent controllers, a combination shelf would not be suitable. Therefore device shelves and two controller shelves would be required. The minimum number of device shelves required to hold the 18 drives of the CAD application would be 6 (See Table 4–2 and plan on the disks being evenly distributed across the two controllers). The 6 device shelves will provide enough additional mounting space to accommodate the 3 drives required for the file services application. The requirement for device shelves and controller shelves would be entered in the Shelves/Enclosures rows of the Configuration Planning worksheet as shown in Example 4–2 (See lightly shaded area).

**Note**

Due to the SCSI bus connections made to the devices shelves and the decision to use a TLZ6L tape loader for backup of the fictional enterprise, one additional shelf is required.



**Example 4–2 Configuration Planning Worksheet—After Adding Shelf Selections**

<b>Data Category—Active (All Users)</b>				
Storage Capacity	10 MB (10 MB)			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures	24-device combination shelf			
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures	6 device shelves		2 controller shelves (shared with file services)	

**Example 4-2 (Cont'd) Configuration Planning Worksheet—After Adding Shelf Selections**


<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures	1 device shelf			

### 4.3.2 Controller Shelves

Figure 4–2 shows a typical controller shelf. Controller shelves are normally used when one of the combination shelves (described in Section 4.3.3) is unsuitable. If controller shelves are selected, one or more device shelves will also be needed for the configuration. The number of controller shelves depends on the following:

- Number of controllers in the configuration
- Controller type
- Configuration of the controllers as independent or redundant.

**Figure 4–2 Typical Controller Shelf**




If you plan independent controllers, you need one controller shelf per controller. If you plan redundant controllers, you need one controller shelf for every redundant pair. Controller shelves mount only in the larger StorageWorks enclosures.

Using the example enterprise, if the storage for the CAD and the file services applications are to be accessed through independent controllers, two controller shelves would be required (The controllers will be used as discussed in Chapter 3 to perform double duty). The requirement for two controller shelves would be entered in the Shelves/Enclosures rows of the Configuration Planning worksheet as shown in Example 4–2 (See medium shaded area).

### 4.3.3 Combination Shelves

Figure 4–3 shows a typical combination shelf mounted in a deskside enclosure. Combination shelves provide mounting space for both storage devices and controllers as well as the required power supplies. Because combination shelves have an integrated backplane with the necessary connections between the controllers (two possible) and the SCSI ports (5 or 6 possible), combination shelves provide the additional benefit of increased device capacity with reduced cabling requirements.

**Figure 4–3 Typical Combination Shelf**



CXO-4268A-MC\_R

Using the example enterprise, if the storage for the relational database and the inventory control applications is through redundant controllers, one 24-device combination shelf would satisfy the requirements. The requirement for a combination shelf would be entered in the Shelves/Enclosures rows of the Configuration Planning worksheet as shown in Example 4–2 (See heavily shaded area).

## 4.4 Selecting Power Components

Selecting power components is determined by the total number and type of shelves used, the types of devices, and whether you desire power redundancy. You require one power supply per device or controller shelf and can install one additional power supply for redundancy. If you plan to use the Combination shelf designed for use in the SW300 (can also be used in the SW800), you require a minimum of 4 power supplies and can install from one to four additional power supplies for redundancy.

Using the example enterprise, redundant power has been specified for all controllers and storage devices. A total of two power supplies per device shelf, two per controller shelf, and five, six, seven, or eight for the combination shelf are required. If eight power supplies for the combination shelf were decided on, the Power Supply rows of the Configuration Planning Worksheet would be filled out as shown in Example 4–3 (See lightly shaded areas).

## 4.5 Selecting Enclosures

Selecting the enclosure depends on the following factors:

- Number of storage devices and controllers
- Available site space
- Available site power

Select the enclosure or enclosures that will hold the number of controller and devices you determined would be required

Using the example enterprise, an SW800 cabinet was selected to hold everything selected. The Shelves/Enclosures rows of the Configuration Planning Worksheet would be filled out as shown in Example 4–3 (See heavily shaded areas)

### Example 4–3 Configuration Planning Worksheet—After Selecting Power Supplies and Enclosures

Data Category—Active (All Users)				
Storage Capacity	10 MB (10 MB)			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				

**Example 4-3 (Cont'd) Configuration Planning Worksheet—After Selecting Power Supplies and Enclosures**

<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration				
Power Supply	8 power supplies			
RAID Arrays				
Storage Devices				
Shelves/Enclosures	24-device combination shelf		SW800	
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration				
Power Supply	12 power supplies (4 for the controller shelves, 8 for the device shelves)			
RAID Arrays				
Storage Devices				
Shelves/Enclosures	6 device shelves	2 controller shelves (shared with file services)		SW800
<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration				
Power Supply	(Shared with CAD)			
RAID Arrays				
Storage Devices				
Shelves/Enclosures	1 device shelf		SW800	

## 4.6 Selecting Accessories


Refer to the *StorageWorks Products Catalog* or the *Systems and Options Catalog* for descriptions of the accessories available for StorageWorks products.

## 4.7 The Configuration Planned For the Fictional Enterprise

Figure 4–4 shows how the configuration planned for the fictional enterprise would appear. Figure 4–5 show the logical configuration planned for the Device shelves that will hold the

storage devices for the CAD and the file services application. (See Appendix A for descriptions of the SCSI bus configuration rules.)

**Figure 4-4 View of SW800 Enclosure Housing the Fictional Configuration**



CABINET FRONT

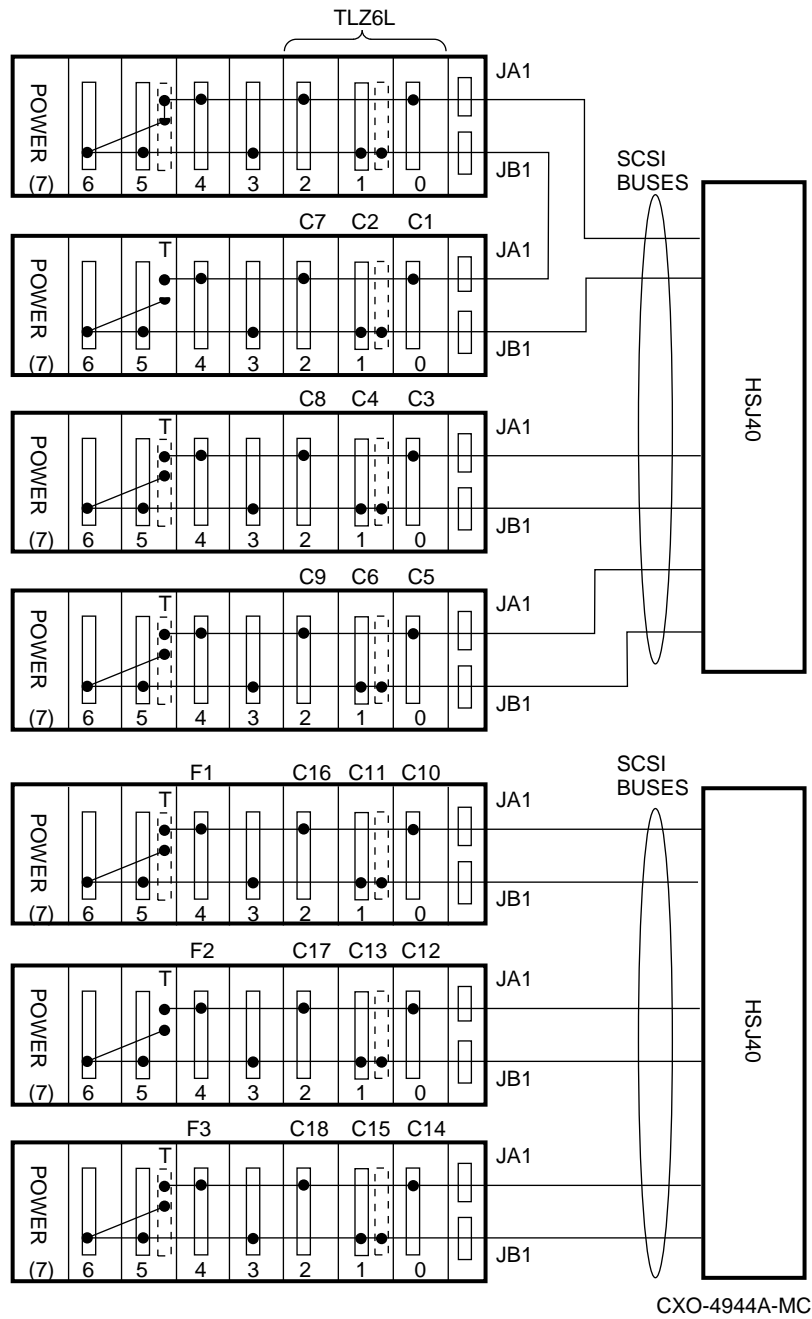
CABINET REAR

CXO-4943A-MC

Where:

- R1 through R16 are the storage devices for the Relational database
- S1 through S4 are the storage devices for the Spareset
- C1 through C18 are the storage devices for the CAD application
- F1 through F3 are the storage devices for the file services application
- The TLZ6L provides backup capabilities for all storage

**Figure 4-5 Logical Configuration of the Device Shelves for the CAD and File Services Applications**





**Example 4–4 Configuration Planning Worksheet—After Completing the Selection Phase**

<b>Data Category—Inactive (Near Line)</b>				
Storage Capacity	1.3GB			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Dormant (Off-Line)</b>				
Storage Capacity	30.25			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	Normal	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				
<b>Data Category—Active (All Users)</b>				
Storage Capacity	10 MB (10 MB)			
I/O Activity	I/O Request Intensive	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration				
Power Supply				
RAID Arrays				
Storage Devices				
Shelves/Enclosures				

**Example 4-4 (Cont'd) Configuration Planning Worksheet—After Completing the Selection Phase**

<b>Data Category—Relational Database and Inventory Control</b>				
Storage Capacity	20 GB (4.4 GB)			
I/O Activity	<b>I/O Request Intensive</b>	<b>Read Intensive</b>	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache; redundant controllers			
Power Supply	Redundant power supplies, 8 power supplies			
RAID Arrays	RAID Level 5; 1 spare			
Storage Devices	RZ28–VA, 16 drives (4 RAIDsets); RZ28–VA 4 drives (spareset)			
Shelves/Enclosures	24-device combination shelf			
<b>Data Category—CAD</b>				
Storage Capacity	15 GB (3.0 GB)			
I/O Activity	I/O Request Intensive	Read Intensive	<b>Data Transfer Intensive</b>	N/A
Data Throughput	N/A			
I/O Response Time	Low	<b>Medium</b>	High	N/A
Data Availability	Normal	<b>Extended</b>	N/A	
Controller Configuration	Read and Write-back cache; Independent controllers			
Power Supply	Redundant power supplies; 12 power supplies (4 for the controller shelves, 8 for the device shelves)			
RAID Arrays	RAID Level 1			
Storage Devices	RZ28–VA, 18 drives (3 RAIDsets)			
Shelves/Enclosures	6 device shelves, 2 controller shelves (shared with file services)			
<b>Data Category—File Services</b>				
Storage Capacity	5 GB (1.0 GB)			
I/O Activity	<b>I/O Request Intensive</b>	Read Intensive	Data Transfer Intensive	N/A
Data Throughput	N/A			
I/O Response Time	<b>Low</b>	Medium	High	N/A
Data Availability	<b>Normal</b>	Extended	N/A	
Controller Configuration	Independent controllers			
Power Supply	Redundant power supplies (shared with CAD)			
RAID Arrays	None			
Storage Devices	RZ28–VA, 3 drives			
Shelves/Enclosures	1 device shelf			

---

## Implementing the StorageWorks Configuration

This chapter provides suggestions to help you plan the implementation phase of the configuration process. The chapter does not discuss specifics since each configuration is customized to satisfy a unique set of requirements.

Implementing a StorageWorks configuration occurs in the following stages:

- Ordering the components
- Preparing the site
- Installing the products
- Integrating the storage system
- Acceptance testing of the storage system
- Populating the storage system with data

### 5.1 Ordering the StorageWorks Components

Ordering the components is accomplished by contacting a Digital Sale Representative. Digital suggests that you discuss the configuration with the Sales Representative. Sales Representatives, because of their knowledge of the entire product line, may be able to suggest a different product mix based on their knowledge of recently released or soon to be released products that may more effectively meet your data storage needs. They will also serve to double check the order numbers and compatibility between the devices selected.

### 5.2 Preparing the Site

Site preparation is the adaptation of the site to meet the space, electrical, and cooling requirements of the storage system before its arrival. In addition, site preparation includes planning for the equipment arrival. Based on the components ordered and the delivery date of the storage system, the site preparation can begin well in advance of equipment arrival.

The Site Preparation Worksheet, shown in Table 5–1, helps you plan the site preparation.

**Table 5–1 Site Preparation Worksheet**

<b>Plan for ...</b>	<b>by ...</b>
equipment arrival	having the necessary personnel and equipment available to accept delivery of the equipment on the scheduled day of arrival.
adequate floor space	checking the specific floor space requirements for each stand alone unit ordered as part of the configuration. Include sufficient space to provide a service area. Insufficient service area could hamper installation and will delay service, should it be required.
proper site power	contracting with a qualified electrician to ensure the power requirements of the equipment are satisfied as well as ensure that existing electrical codes are obeyed.
required cooling	ensuring the site cooling capacity is sufficient to handle the additional load created by the storage system. As with any electronic equipment, StorageWorks products generate heat when operating. The heat must be removed for reliable operation. The specifications for the enclosure and power system specify the heat generated (typically in BTUs).

### 5.3 Installing the Equipment

The StorageWorks system is a custom configuration. It may arrive as one or more enclosures populated with devices specified when ordering, or it may arrive as individual building blocks that require installation.

If the storage system arrives as fully assembled enclosures, they may be able to be placed at the desired floor location, plugged in, powered on, and integrated into the existing system.

If the storage system arrives as individual building blocks, the included documentation details the installation procedures for the individual building blocks. When assembling the storage system, follow the general configuration rules for shelves and for the SCSI bus. Appendix **Error! Not a valid filename.** provides the general configuration rules.

### 5.4 Integrating the Storage System

Integrating the storage system is the process of merging the new storage system with the existing system. Depending on the existing system and the intent of the new storage system, the process can range from simple to very complex. For example, if the new storage system is intended to augment an existing storage system, the process can be quite simple and accomplished with little impact to the normal operations of the host system. On the other hand, if the new storage system is intended to replace all or part of an existing storage system, the process becomes very complex and will probably impact the normal operations.

---

**Note**

---

Before integrating the storage system make sure any existing storage has been backed up and stored in a safe place.

---

Some general questions that must be answered in advance are included in the pre-integration checklist shown in Table 5–2.

**Table 5–2 Pre-Integration Checklist**

Question	Yes	No
Will regeneration of the system software be required?		
Will applications require shutdown while integrating the storage system?		
Will the storage capacities provided by the new storage system exceed the limits of the operating system?		
Will data be migrated between the new and any existing storage systems?		

Any question that has a yes response suggests that you formulate a plan to address the potential impact to normal operations. The following paragraphs provide examples of what may be required for a yes response to each question in the Pre-Integration Checklist.

- If integration requires a new system generation, the entire system becomes unavailable until the system generation completes. All information required to perform a complete system generation should be on hand and ready for use. Schedule the system regeneration during a slack time in normal operations.
- If applications require shutdown during integration of the storage system, what degree of shutdown is required? The integration may be able to proceed while the applications execute with reduced functionality
- If the increased storage capacities provided by the new storage system are beyond the limits supported by the operating system, partitioning of the storage devices may be necessary. Device partitioning is the process of dividing a logical device into two or more parts.
- If data will be migrated between the new storage system and an existing storage system, have all system performance issues with the existing storage system been resolved? Migration of data from one system to another may uncover performance issues that were previously hidden and make later identification of the problem source more difficult than necessary.

## 5.5 Acceptance Testing of the Storage System

Acceptance testing is a predetermined series of exercises intended to stress the equipment and identify faults or bottlenecks. Acceptance testing can be an extended series of diagnostic tests or a more thorough acceptance test involving running real applications and test data. The level of acceptance testing performed and the criterion that determines success or failure should be established in advance of storing any data that has real value to the enterprise on the new storage system.

## 5.6 Populating the Storage System With Data

Populating the storage system is the process of moving data to it. The data may be created by executing applications and stored directly to the storage system, or the data may be migrated to the new storage system from an existing storage system. In the latter case, plans must be made to avoid impacting normal system operation. If applications will have to be shut down or provide less than normal functionality to users during the migration, it may be beneficial to schedule the migration to a time period when the applications are not in demand. For example, migration of the data required for a payroll application that runs every other week

would best be migrated during the weekend after the payroll application had been run. This allows up to two weeks to fix unexpected problems that surface after the data migration.

---

## General Configuration Rules

This appendix contains the general configuration rules that apply to configuration of a StorageWorks system. The rules are grouped as follows:

- Controller configuration rules
- Shelf configuration rules
- Power system configuration rules
- SCSI bus configuration rules

### A.1 Controller Configuration Rules

The following general configuration rules apply to StorageWorks controllers:

- Two solid state disks may consume the I/O capability of a single StorageWorks family controller. Therefore, when solid state disks are attached, all additional devices attached to that controller should be low-I/O utilization devices (bulk storage, tapes, etc.).
- Different types of devices may require their own ports on the HS family of controllers. The devices listed on a single line below can be grouped together on the same SCSI port:
  - Magnetic and solid state disks
  - Optical Disks, CD ROM, tapes, and tape loaders
  - Tape libraries
  - Optical jukeboxes

### A.2 SBB Shelf Configuration Rules

Shelf configuration rules are based on the number of controller ports and the type of storage devices to be configured. The following sections describe the configuration rules for:

- 3.5-Inch SBB configurations on 3-port controllers
- 5.25-Inch SBB configurations on 3-port controllers
- 3.5-Inch SBB configurations on 6-port controllers
- 5.25-Inch SBB configurations on 6-port controllers

## A.2.1 3.5-Inch SBB Configurations on 3-Port Controllers

Table A–1 describes how to configure the maximum number of SBB shelves to properly allocate 3.5-inch storage devices across 3–port controllers such as the CI-to-SCSI and DSSI-to-SCSI controllers.

**Table A–1 3.5-Inch SBB Configurations on 3-Port Controllers**

Devices	Shelves	Configure as	Empty positions	Ports used
1 or 2	1	(1) 2X3T*	5 or 4	1 or 2
3 to 12	2	(1) 2X3T (1) 1X6T†	9 to 0	3
13 to 18	3	(3) 1X6T	5 to 0	3
* 2X3T means one shelf with two split SCSI–2 connections, separately terminated in each shelf. † 1X6T means one shelf with a single SCSI–2 connection terminated in the shelf.				

## A.2.2 5.25-Inch SBB Configurations on 3-Port Controllers

Table A–2 describes how to configure the maximum number of SBB shelves to properly allocate 5.25-inch storage devices across 3–port controllers such as the CI-to-SCSI and DSSI-to-SCSI controllers.

**Table A–2 5.25-Inch SBB Configurations on 3-Port Controllers**

Devices	Shelves	Configure as	Empty positions	Ports used
1 or 2	1	(1) 2X3T*	1 or 0	1 or 2
3 or 4	2	(1) 2X3T (1) 1X6T†	1 or 0	3
5 or 6	3	(3) 1X6T	1 or 0	3
7 or 8	4	(3) 1X6T (1) 1X6J‡	1 or 0	3
9 to 10	5	(3) 1X6T (2) 1X6J	1 or 0	3
11 or 12	6	(3) 1X6J	1 or 0	3
* 2X3T means one shelf with two split SCSI–2 connections, separately terminated in each shelf. † 1X6T means one shelf with a single SCSI–2 connection terminated in the shelf. ‡ 1X6J means one shelf, not terminated, that is jumpered to the next shelf				



### A.2.3 3.5-Inch SBB Configurations on 6-Port Controllers

Table A–3 describes how to configure the maximum number of SBB shelves to properly allocate 3.5-inch storage devices across 6-port controllers such as the CI-to-SCSI and differential-to-single-ended SCSI controllers.

**Table A–3 3.5-Inch SBB Configurations on 6-Port Controllers**

Devices	Shelves	Configure as	Empty positions	Ports used
1 or 2	1	(1) 2X3T*	5 or 4	1 or 2
3 or 4	2	(2) 2X3T	9 or 8	3 or 4
5 to 18	3	(3) 2X3T	13 to 0	5 or 6
19 to 24	4	(2) 2X3T (2) 1X6T†	5 to 0	6
25 to 30	5	(1) 2X3T (4) 1X6T	5 to 0	6
31 to 36	6	(6) 1X6T	5 to 0	6

\* 2X3T means one shelf with two split SCSI–2 connections, separately terminated in each shelf.  
† 1X6T means one shelf with a single SCSI–2 connection terminated in the shelf.

### A.2.4 5.25-Inch SBB Configurations on 6-Port Controllers

Table A–4 describes how to configure the maximum number of SBB shelves to properly allocate 5.25-inch storage devices across 6-port controllers such as the CI-to-SCSI and differential-to-single-ended SCSI controllers.

**Table A–4 5.25-Inch SBB Configurations on 6-Port Controllers**

Devices	Shelves	Configure as	Empty positions	Ports used
1 or 2	1	(1) 2X3T*	1 or 0	1 or 2
3 or 4	2	(2) 2X3T	1 or 0	3 or 4
5 or 6	3	(3) 2X3T	1 or 0	5 or 6
7 or 8	4	(2) 1X6T† (2) 2X3T	1 or 0	6
9 or 10	5	(4) 1X6T (1) 2X3T	1 or 0	6
11 or 12	6	(6) 1X6T	1 or 0	6
13 or 14	7	(6) 1X6T (1) 1X6J‡	1 or 0	6
15 or 16	8	(6) 1X6T (2) 1X6J	1 or 0	6
17 or 18	9	(6) 1X6T (3) 1X6J	1 or 0	6
19 or 20	10	(6) 1X6T (4) 1X6J	1 or 0	6
21 or 22	11	(6) 1X6T (5) 1X6J	1 or 0	6
23 to 24	12	(6) 1X6T (6) 1X6J	1 or 0	6

\* 2X3T means one shelf with two split SCSI–2 connections, separately terminated in each shelf.  
† 1X6T means one shelf with a single SCSI–2 connection terminated in the shelf.  
‡ 1X6J means one shelf, not terminated, that is jumpered to the next shelf

### A.3 Power System Configuration Rules

The following general configuration rules apply to StorageWorks power system:

- Each shelf requires either an ac or dc input primary power supply.
- The shelf power supply must be mounted in slot 7 (power slot) of the shelf.
- An optional redundant power supply or an optional battery backup unit (BBU) can be mounted in slot 6.

### A.4 SCSI Bus Configuration Rules

The StorageWorks data bus is a single-ended, small computer system interface (SCSI) bus and is available as an 8-bit or a 16-bit wide data path. The following general configuration rules apply to StorageWorks SCSI buses:

- The maximum number of storage devices on a single 8-bit wide SCSI-2 bus is seven; the maximum number of storage devices on a 16-bit wide SCSI-2 bus is 15.
- The total length of either the 8-bit or the 16-bit wide single-ended SCSI-2 bus is the distance between the two bus terminators.
- The total length of either the 8-bit or 16-bit wide single-ended SCSI-2 bus operating at 10 MB/s cannot exceed 3.0 meters (9.8 feet).
- The total length of either the 8-bit or the 16-bit wide single-ended SCSI-2 bus operating at 5 MB/s cannot exceed 6.0 meters (19.7 feet).
- The rate at which data is transferred over either the 8-bit or the 16-bit single ended SCSI-2 bus is primarily determined by the capabilities of the host SCSI adapter or controller. A controller is rated as either fast or slow based on the maximum rate at which it can transfer data. Prior to transferring the data, the controller determines the maximum rate at which each device can accept data. If this rate is lower than the controller's maximum rate, the controller uses the lower rate for the transfer.

This capability enables the reliable transfer of data to both slow and fast devices on the same bus, provided the bus is designed as a fast bus with a maximum bus length of 3.0 meters (9.8 feet).

---

#### Note

---

When the adapter can transmit at the fast rate and there is at least one fast device on the bus, the SCSI bus *must be designed* as a fast bus (maximum length 3.0 meters [9.8 feet]).

---

- The total length of a fast-wide, 16-bit, differential SCSI bus cannot exceed 25 meters (82 feet).
- Both a single-ended and a differential SCSI bus can be extended using SCSI signal converters. A maximum of two SCSI signal converters, including one installed in the shelf, such as a DWZZA-VA, can be used on any one bus.
- A single, fast SCSI bus cannot be used for more than two shelves.
- All SCSI buses require termination at both ends of the bus.
- The single-ended SCSI buses are terminated on the last shelf on the bus.

---

**Note**

---

SCSI buses must be terminated at both ends for proper operation. SBBs do not terminate a bus; a BA35X-MB terminator board must be installed on the shelf backplane.

---

- The rules for installing SCSI bus terminators and jumpers are shelf specific.
- Multiple SCSI buses on a shelf can be connected together using a jumper board (BA35X-MC). For example, two buses can be connected into one; three buses can be connected into two, and so on.
- The target ID addresses for hosts and controllers are normally as follows:
  - First (single) host or controller: ID 7
  - Second (dual) host or controller: ID 6
- Device address 7 is normally reserved for the host computer address.
- The remaining seven addresses (0 through 6, or 0 through 15) are used for the individual storage device addresses.

The maximum length of the fast and slow buses, including all cables and shelf buses, is shown in Table A-5. All bus lengths are rounded off to the nearest tenth of a unit.

**Table A-5 SCSI Bus Parameters**

Bus Type	Transfer Rate	Meters	Feet
single-ended (8-bit or 16-bit)	5 MB/s	6.0	19.7
single-ended (8-bit or 16-bit)	10 MB/s	3.0	9.8
differential (16-bit)	20 MB/s	25.0	82.0

#### A.4.1 Terminating SCSI Buses

The shelf end of the SCSI bus is terminated with a terminator board mounted on the shelf backplane. See the controller documentation for termination options.

---

**Note**

---

All StorageWorks SBB shelf SCSI buses are terminated on the shelf backplane. There is no SCSI bus termination on any SBB device.

---

#### A.4.2 Computing SCSI Bus Lengths

Operation of an 8-bit, single-ended, SCSI bus is affected by the total bus length (that is, the total distance between the two bus terminators). Failure to accurately determine the total SCSI bus length can cause signal degradation to the point where the bus is non-operational. Practical experience, especially when dealing with a fast SCSI bus (10 MB/s transfer rate), has shown that the limitations stated in the previous section titled StorageWorks Data Bus Description are applicable for most bus operations.

Use the procedures in the following sections to determine the SCSI bus lengths for either single or multiple SBB shelves.

### A.4.3 Single StorageWorks Shelf SCSI Bus Lengths

As shown in Figure A-1, the SCSI bus for a single shelf is composed of the following segments:

- **Segment A**—The distance from the controller terminator (T1) to the SCSI bus connector (BC).
- **Segment B**—The length of the SCSI cable from the controller shelf to the SBB shelf (either 1 meter [3.3 feet] or 2 meters [6.5 feet]).
- **Segment C**—For a single SBB shelf this is the distance from the SBB shelf input connector (IC) to the shelf SCSI terminator (T2).

Use the information in Table A-6 and the following formula to calculate the length of a SCSI bus:

$$\text{Total Bus Length} = A+B+C+I(n)$$

---


**Note**

---

The last term in the equation  $I(n)$  allows for the SCSI bus that is internal to the storage device. A value of 1 foot (I) multiplied by the number of devices (n) provides an adequate allowance.

---


**Figure A-1 SCSI Bus Segments — Single SBB Shelf**



CXO-3941A-MC

**Table A-6 Single SBB Shelf SCSI Bus Lengths**


Shelf	Bus	Input	Slot 6	Slot 1	Meters	Feet
<b>Single Shelf—Single Bus</b>						
1	1	JA1	Jumper	None	1.0	3.3



CXO-4513A-MC

**Single Shelf—Two Buses**

1	1	JA1	Terminator	Jumper*	0.4	1.4
1	2	JB1	Terminator	Jumper*	0.6	2.0



CXO-4514A-MC

\* The jumper is placed here for safekeeping and has no effect on the SCSI bus.

### A.4.4 Multiple StorageWorks Shelves SCSI Bus Lengths

As shown in Figure A-2, the SCSI bus for multiple SBB shelves on a single SCSI bus is composed of the following segments:

- **Segment A**—The distance from the controller terminator to the shelf backplane connector (port) and from there to the SCSI bus connector
- **Segment B**—The length of the BN21H-series cable from the controller shelf to the SBB shelf (either 1 meter [3.3 feet] or 2 meters [6.5 feet])
- **Segment C**—The distance from the first or second SBB shelf input connector through the jumper (J) to the shelf output connector
- **Segment D**—The length of the SCSI cable (0.5 meters or 1.6 feet) that connects the SBB shelves
- **Segment E**—The distance from the last shelf input connector to the bus terminator on the same shelf

Use the information in Table A-7 and the following formula to calculate the length of a SCSI bus:

$$Total\ Bus\ Length = A+B+C+D+E+I(n)$$

**Note**

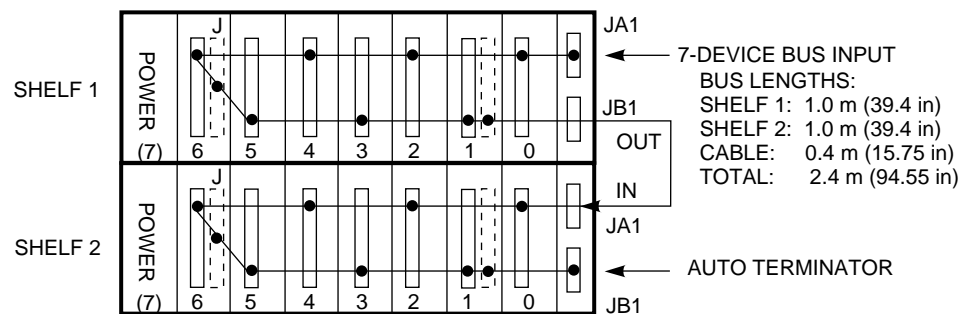
The last term in the equation I(n) allows for the SCSI bus that is internal to the storage device. A value of 1 foot (I) multiplied by the number of devices (n) provides an adequate allowance.

**Table A-7 Multiple SBB Shelves SCSI Bus Lengths**

Shelf	Bus	Input	Slot 6	Slot 1	Meters*	Feet*
<b>Two Adjacent Shelves—Single Bus</b>						
1	1	JA1	Jumper	None†	1.0‡	3.3
2	1	JA1	Jumper	None	1.0	3.3
			BN21L-0E Cable		0.5	1.6
			Total Bus Length		2.5	8.2


‡ From the input connector (JA1) to the output connector (JB1).

† Do not install a terminator here.



CXO-4515A-MC

**Figure A-2 SCSI Bus Segments — Multiple SBB Shelves**



\*WHEN THERE ARE MORE THAN TWO SHELVES, REPEAT SEGMENTS C AND D FOR ALL BUT THE LAST SHELF.

CXO-3942A-MC





---

## Raid Overview

RAID (Redundant Array of Independent Disks) is an array of disks working together under the control of Array Management Software to provide one or more virtual disks that provide greater data availability and better disk performance than can be attained by a simple collection of disks. Raid distributes the data across multiple disks, provides, except in the case of Striping (RAID Level 0), data redundancy, and improves I/O performance for the following types of I/O:

- Data transfer-intensive—Reading or writing large amounts of data per unit time, usually in a single stream. Applications such as computer assisted design, image processing, data collection and reduction, and sequential file processing tend to be data transfer-intensive.
- I/O request-intensive—Executing large numbers of I/O requests per unit time. Applications such as interactive transaction processing, office automation, and providing general I/O or file services for general timesharing tend to be I/O request-intensive.

Digital supports RAID Level 0, Mirroring (RAID Level 1), Parity RAID (RAID Level 5), and Striped Mirrorsets (RAID Level 0+1). The following sections present a brief overview of the RAID levels supported by Digital and the type of applications for which they are best suited. See other documentation, such as the RAIDbook published by the RAID Advisory board for details about the implementation of each RAID level.

---

### Note

---

If you choose a RAID solution, there will be an added cost ranging from one additional disk per array to a complete duplication of the RAID array.

---

## B.1 Striping (RAID Level 0)

RAID Level 0 is a performance oriented mapping technique. Because it does not provide data redundancy (the only RAID level that does not), it provides no provision against data loss due to member disk failures. The failure of a single member of a RAID Level 0 array is equivalent to a failure of the entire array.

RAID Level 0 is particularly well suited for the following:

- Storing program image libraries or run-time libraries for rapid loading. These libraries are normally read only. They can be backed up on highly reliable media from which they can be re-created in the event of a failure.
- Storing large tables or other structures of read-only data for rapid application access. Like program images, these data can be backed up on highly reliable media, from which they can be re-created in the event of a failure.
- Collecting data from external sources at very high data transfer rates. This works best where the data is reproducible by repeating the process or experiment which produces it. In some cases, the acquisition rate and volume of data may be so high that there is no alternative to RAID Level 0.

RAID Level 0 is not particularly well suited for the following:

- A storage solution for data that cannot be easily reproduced or for data that must be available for system operation.
- Applications that make requests for small amounts of sequentially located data.
- Applications that make synchronous random requests for small amounts of data.

If RAID Level 0 performance and high data availability are required, one solution is to combine RAID Level 0 and RAID Level 1 (RAID Level 0&1).

## **B.2 Mirroring (RAID Level 1)**

RAID level 1 provides redundancy of data by shadowing the entire virtual disk at least once. RAID Level 1 provides an available copy of any user visible data, even if a member disk fails. In addition, a RAID Level 1 configuration can improve the performance of read intensive applications because the Array Management Software chooses which member disk should handle the read request. Some implementations choose alternately or randomly among member disks for a simple form of load balancing. More sophisticated designs select the least busy member to better balance the I/O load.

## B.3 RAID Level 3

Although RAID Level 3 is not supported by StorageWorks, an explanation of it follows for those who are interested. RAID Level 3 adds a parity disk to a striped array (RAID Level 0), permitting user data to be regenerated in the event of a single disk failure. In general RAID Level 3 has the following characteristics:

RAID Level 3 is particularly well suited for the following:

- Applications that require high data transfer capacity
- Applications that operate on large data objects such as CAD files, graphical images, seismic or telemetered data streams.
- Non-interactive applications that process large files sequentially.

RAID Level 3 is not particularly well suited for the following:

- Transaction processing
- I/O request rate intensive

However, the use of some other technology, such as caching, can make RAID Level 3 an alternative configuration choice for those application types as well.

---

### Note

---

Because every member of a RAID level 3 array participates in every I/O operation, the array's aggregate member disk data transfer capacity must be continuous along the system's entire data transfer path.

---

## B.4 Parity RAID (RAID Level 5)

RAID Level 5 stripes data and parity across its member disks and has the following performance characteristics:

- The aggregate I/O performance of a RAID Level 5 array depends on the percentage of reads and writes in the I/O load.
- RAID Level 5 arrays are usually more suitable for I/O request-intensive applications than for data transfer-intensive applications.
- Application transactions that include large numbers of writes will exhibit poorer response when they write to a RAID Level 5 array than when they write to individual disks whose performance is equivalent to that of the array's member disks.
- RAID Level 5 arrays to which a high percentage of write requests are directed will saturate at a lower I/O request rate than individual disks whose performance is equivalent to that of the array members.

RAID Level 5 is particularly well suited for the following:

- Inquiry-type transaction processing
- Group office automation
- On-line customer service departments

RAID Level 5 is not particularly well suited for the following:

- High-speed data collection
- Database applications in which fields are continually updated

---

**Note**

---

When Raid Level 5 technology is combined with write-back cache to improve its write performance, the resulting array can be used in any application where general purpose disks would be suitable.

---

## B.5 Striped Mirrorsets (RAID Level 0 & 1)

RAID Level 0&1 combines RAID level 0 and Raid Level 1 in a single array that provides improved I/O performance through RAID Level 0 and improved data reliability through RAID Level 1. The result is storage with very high I/O performance and high data reliability, however at a relatively high hardware cost.

Raid Level 0&1 performance is due to the following:

- The RAID Level 0 array balances the I/O load. For data transfer-intensive applications, load balancing occurs because concurrent I/O requests for relatively small amounts of data tend to be distributed evenly across the member disks. For I/O request-intensive applications, load balancing occurs because each application I/O request is broken down into multiple requests which are executed concurrently by RAID Level 0 array members.
- When the I/O load is high, I/O queues for the RAID Level 1 member disks of the RAID Level 0 array may build up. When this occurs, RAID level 1 array management software can produce a secondary level of load balancing by choosing the least busy member disk to satisfy each read request.

These effects can significantly improve performance for I/O request-intensive applications which meet the basic requirements for I/O load balancing:

- I/O request-intensive applications with I/O loads of asynchronous I/O requests
- Data transfer-intensive applications which make large double buffered requests

---

### Note

---

RAID Level 0&1 is an excellent choice for any data whose value justifies placing it on shadowed storage.

---

## B.6 Summary Comparison of RAID Levels

The following table provide a quick comparison chart of the most common RAID levels.

RAID Level	Common Name	Description	Disks Required	Data Reliability	Data Transfer Capacity	Maximum I/O Rate
0	Disk Striping	Data is distributed across the disks in the array No redundant information provided.	N	Lower than a single disk	Very high	Very high for both read and write
1	Shadowing	All data is replicated on N separate disks. N is almost always 2.	2N, 3N, etc.	Higher than RAID Level 2,3,4, or 5: lower than RAID Level 6	Higher than single disk for read; similar to single disk for write	Up to twice that of a single disk for read; similar to a single disk for write.
3	RAID 3, Parallel Transfer Disks with Parity	Each data sector is subdivided and distributed across all data disks. Redundant information normally stored on a dedicated parity disk	N+1	Much higher than a single disk; comparable to RAID 2, 3, or 5	Highest of all listed alternatives	Similar to twice that of a single disk
5	RAID 5, RAID	Data sectors are distributed as with disk striping; redundant information is interspersed with user data	N+1	Much higher than a single disk; comparable to RAID 2, 3, or 4	Similar to disk striping for read; significantly lower than single disk for write	Similar to disk striping for read; significantly lower than single disk for write





---

## Migration Planning— An Overview

---

### Note

---

A detailed discussion of migration planning is beyond the scope of this Configuration Guide. The information provided in this appendix touches only the surface of a very complex subject. Digital recommends that you contact Digital Sales for assistance in formulating a migration plan as part of your configuration analysis

---

Migration planning is preparing a strategy for replacing one storage system with another, transferring the data of the storage system being replaced to the replacement storage system, and maintaining the same or better levels of performance and data availability for all users and applications that access data on the new storage system. Migration planning has two objectives:

- Identify the current system in terms of its I/O requirements and performance
- Determine the types and numbers of replacement storage products

The first objective is met by a complete system analysis. The second objective is met by using the results of the analysis to design and configure a subsystem that meets or exceeds the requirements and performance demanded of the original storage subsystem.

### C.1 Identify the Current System In Terms of Its I/O Requirements and Performance

The analysis is accomplished in steps as follows:

- Identify any system level or storage subsystem performance problems
- Identify the applications, software, and uses of the current system
- Determine the expected storage capacity and performance requirements
- Identify which part of the storage subsystem will be migrated
- Identify the data backup methodologies and products used
- Identify the system and storage/controller configuration
- Identify the methods used to provide device availability
- Determine if disk storage is fragmented

### **C.1.1 Identify Any System Level or Storage Subsystem Performance Problems**

The result of migrating to another storage subsystem can worsen existing system level, or I/O subsystem problems. Any problems with the current system must be identified and resolved prior to an analysis of the storage system performance. If problems are present in the current system, those problems will prevent an accurate measurement of the real I/O activity that must be planned for.

The potential I/O performance of a system cannot be measured accurately if system problems such as insufficient memory or CPU resources are present. For example, insufficient memory causes excessive paging and swapping. The I/O required for paging and swapping subtracts from the I/O that would otherwise be available for the use of executing applications. Insufficient CPU resources cause an apparent decrease in the I/O requirements of a system. As jobs wait for CPU resources to become available, the jobs are not requesting the I/O they would normally. Both of these situations would cause an incorrect measurement of the system potential for satisfying application I/O demands and result in an incorrect analysis of the system I/O requirements.

If system problems such as these are identified, they should be resolved before formulating a migration plan. After system problems such as these are resolved, the resultant increase in I/O by applications may uncover real I/O bottlenecks which can be resolved as part of the migration plan.

### **C.1.2 Identify the Applications, Software, and Uses of the Current System**

Different applications and layered software products place different demands on the storage system depending on their purpose. For example, high throughput is critical for satellite data acquisition, whereas, most transaction processing applications require fast response times. Knowing which applications use what resources and if the applications run concurrently or at different times is important in determining placement of the required files during migration.

### **C.1.3 Determine the Expected Storage Capacity and Performance Requirements**

Current and future capacity growth plans, when factored into the migration plan, may provide a cost savings over the long term. For example, the anticipated growth of the storage subsystem may become a determining factor in the choice of storage devices or enclosures. These decisions could then impact the site planning that will be required as part of the migration plan.

### **C.1.4 Identify Which Part of the Storage System Will Be Migrated**

The migration plan can focus on only the devices that will be replaced.

### **C.1.5 Identify the Data Backup Methodologies and Products Used**

Increased capacity can cause difficulties if the backup device capacity or methodology is not adequate.

### **C.1.6 Identify the System and Storage/Controller Configuration**

The typical result of a storage subsystem migration is the use of the newest products as replacements. Frequently this requires microcode and/or hardware updates to the controllers or other system components. By knowing about configurations and compatibility early in the process, arrangements can be made for any further upgrades before they become critical at replacement device installation time. Also, different system/controller/device configurations can affect I/O performance positively or negatively and must be factored into the analysis.

If disk devices are configured in a shadow set, the migration plan should propose enough additional disks to maintain the shadow set or an alternative to the shadow set must be proposed.

If stripe sets are being used determine which disks comprise the stripe sets and how they are being used in the current system. Determine if the number of members in the strip set are to remain the same or change. If the number remains the same, the use of large capacity disks (relative to the existing disk capacity) may not be an option.

### **C.1.7 Identify the Methods Used to Provide Device Availability**

Migrating to a fewer number of higher capacity disks could affect data availability schemes such as device failover.

### **C.1.8 Determine if the Disk Storage is Fragmented**

Disk fragmentation causes additional I/O requests that must be accounted for in the I/O performance analysis. Defragmenting the disks should be a part of the migration plan.

## C.2 Determine the Types and Numbers of Replacement Storage Products

Based on the information provided by the analysis, begin determining the types and number of storage products for the new configuration. The most important information on which the determinations should be made are the following:

- The I/O demands of users and applications running on the system
- What files, users, and applications are associated with each other

The I/O demands of users and applications must be met with the same or better level of response as the previous storage system. In order to accomplish this the association of files, users, and applications to one another must be known so the files can be positioned across the replacement products in a way that does not create I/O problems.

For example, assume the original storage system has two disks configured as general user disks and each handled 30 I/Os per second. If, after migration, the files on the two disks end up on one higher capacity replacement disk, the users will experience a significant degradation in throughput as the replacement disk will have to achieve 60 I/Os per second.

The relationship of files to one another and to applications must be carefully considered when placing files on the replacement system. By incorrectly placing files, the benefits of the new system can be negated and the performance of the replacement system can be worse than the replaced system.

For example, assume a payroll application requires separate files containing payroll and employee benefits information, and those files are on separate disks in the original storage system. If, after migration, the files end up on one of fewer, but larger capacity, replacement disks, the payroll applications will run slower. The application runs slower because of I/O contention problems created by file placement.

The replacement of one storage system with another based on capacity and lower initial costs often results in fewer, higher capacity replacement storage devices. A reduction in the “spindle count” can easily degrade the I/O performance of the storage system. Analyzing the I/O requirements can reveal that more than a device replacement strategy is required. A comprehensive solution may require different storage device technologies (for example, optical devices, tape libraries, or solid state disks, and RAID), software products such as file optimizers, additional memory, or additional CPU power to name a few.

---

### Note

---

I/O performance is a relative measurement of how well the I/O system responds to application demands for data transfers. If the demands are responded to immediately, the I/O performance is very high. As more time is required to respond to the demand, the I/O performance decreases.

---

---

# Glossary

**ac distribution**

The method of distributing ac power in a cabinet.

**ac power supply**

A power supply designed to produce dc power from an ac input.

**adapter**

A device that converts the protocol and hardware interface of one bus type into that of another without changing the functionality of the bus. See signal converter.

**American National Standards Institute**

See ANSI.

**ANSI**

American National Standards Institute. An organization that develops and publishes electronic and mechanical standards.

**array controller**

A hardware/software device that facilitates communications between a host and one or more devices organized in an array. The HSJ40 controller is an array controller.

**BA350-Mx controller shelf**

The StorageWorks controller shelf used for HS-family controller modules, cache modules, and shelf power units.

**BA350-Sx SBB shelf**

A StorageWorks shelf used for only power units and SBBs.

**battery backup unit**

See BBU

**BBU**

StorageWorks battery backup unit that extends power availability after the loss of primary ac power or a power supply to protect against the corruption or loss of data. blower An airflow device mounted in a StorageWorks shelf.

**cable distribution unit**

See CDU.

**carrier**

A standard, StorageWorks shelf-compatible, plastic shell into which a device can be installed. Sometimes called SBB carrier.

**CDU**

Cable distribution unit. The power entry device for StorageWorks center cabinets. The unit provides the connections necessary to distribute ac power to cabinet shelves and fans.

**channel**

The RAID Array 110 subsystem term for a SCSI bus.

**CI bus**

Digital's computer interconnect bus using two serial paths, each with a transfer rate of 70 Mb/s (8.75 MB/s).

**cold swap**

A method of device replacement that requires that power be removed from all shelves in a cabinet. This method is used when conditions preclude the use of the warm swap or hot swap methods.

**configure-to-order**

See CTO.

**controller**

A hardware/software device that facilitates communications between a host and one or more devices.

**controller shelf**

A StorageWorks shelf designed to contain controller and cache memory modules.

**CTO**

Configure-to-order. The customer-specified options are installed at the factory.

**DAT**

Digital Audio Tape. A format for recording digital data on a cartridge tape.

**data center cabinet**

A generic reference to the large cabinets, such as the SW800 series, in which StorageWorks components can be mounted.

**data channel**

An HSC interface between the disk drive or tape formatters and HSC data and control memory.

**deskside expansion unit**

A pedestal enclosure designed to house one or two StorageWorks shelves in a vertical position.

**DIGITAL Storage System Interconnect**

See DSSI

**DSSI**

Digital's storage system interconnect bus with an 8-bit data transfer rate of 4-5 MB/s.

**dual cabinet power configuration**

A cabinet ac power configuration in which two ac sources and two ac power supplies are used to supply dc power to the cabinet's SBB shelves.

**dual porting (or dual access)**

The ability of a disk or tape drive to be accessed by two controllers. All DSA drives have a standard dual-port feature. DSA drives can be online to only one controller at a time. However, they are able to disconnect themselves from a failed controller (or be disconnected by a failing controller) and become available for continued service through the other controller.

**dual-redundant configuration**

A controller configuration consisting of a primary and backup controller in one controller shelf. If the primary controller fails, the backup controller assumes control over the failing controller's devices.

**dual shelf power configuration**

A cabinet ac power configuration in which one ac source and two ac power supplies are used to supply dc power to the cabinet's SBB shelves.

**electromagnetic interference**

See EMI

**electrostatic discharge**

See ESD

**EMI**

Electromagnetic interference. The impairment of a signal by an electromagnetic disturbance.

**ESD**

Electrostatic discharge. The discharge of a potentially harmful static electric voltage as a result of improper grounding.

**factory-configured**

These systems are configured at the factory and shipped to the customer "as is." Customer-specified options are shipped with the system for on-site installation.

**failover**

Failover is the software process that takes place when one controller fails in a dual-redundant configuration and the other controller takes over the devices of the failed controller and services them to the host until or if the failed controller comes back or is replaced.

**fan**

An airflow device mounted in a StorageWorks cabinet.

**fast, differential SCSI**

See FD SCSI. fast, wide, differential SCSI See FWD SCSI.

**FD SCSI**

The fast, differential SCSI bus with an 8-bit data transfer rate of 10 MB/s. See also FWD SCSI and SCSI.

**field configured**

Customer-specified options are installed on-site. The options can be ordered at any time for field upgrades.

**full-height device**

A single device that occupies an entire 5.25 inch SBB carrier. StorageWorks full-height devices have an order number suffix of "-VA".

**FWD SCSI**

The fast, wide, differential SCSI bus with a 16-bit data transfer rate of up to 20 MB/s. See also FD SCSI and SCSI.

**half-height device**

A device that occupies half of a 5.25 inch SBB carrier. Two half-height devices can be mounted in a 5.25 inch SBB carrier. The first half-height device is normally mounted in the lower part of the carrier. The second device is normally mounted in the upper part of the carrier.

**Hierarchical Storage Controller**

See HSC.

**host**

The primary or controlling computer to which a storage subsystem is attached.

**hot swap**

A method of device replacement whereby the complete system remains on line and active during device removal and reinstallation. The device being removed or reinstalled is the only device that cannot perform operations during this process.

**HSC**

Hierarchical Storage Controller. An intelligent mass storage server used on the CI bus. Capable of supporting a total of eight disk and/or tape data channels, the HSC is part of the System Interconnect Architecture and Digital Storage Architecture. By performing as an I/O manager, the HSC can be classified as an I/O server, removing the burden of I/O management from the CPU.

**qualified device**

A device that has been fully tested in all appropriate StorageWorks hardware and software configurations, and is in complete compliance with Digital and country-specific standards (for example, FCC and TÜV).

**radio frequency interference**

See RFI.

**RAID**

Redundant array of independent disks. A set of storage techniques devised to increase the performance and availability of a storage subsystem.

**redundant array of independent disks**

See RAID.

**RFI**

Radio frequency interference. The impairment of a signal by an unwanted radio signal or radio disturbance.

**SBB**

StorageWorks building block. A device housed in a standard StorageWorks SBB carrier. An SBB has a standard physical and electrical interface that is compatible with those of StorageWorks shelves and enclosures.

**SBB shelf**

A StorageWorks shelf, such as the BA350-Sx, designed to house plug-in SBB modules.

**SCSI**

Small Computer System Interface. An ANSI interface defining the physical and electrical parameters of a parallel I/O bus used to connect hosts to a maximum of seven devices. The StorageWorks device interface is implemented according to the SCSI-2 standard, allowing the synchronous transfer of 8-bit data at rates of up to 10 MB/s.

**shelf brackets**

Sheet metal components designed to attach and position StorageWorks shelves in their associated enclosures.

**signal converter**

A device that converts the protocol and hardware interface of one bus type into that of another without changing the functionality of the bus. See adapter.

**Small Computer System Interface**

See SCSI.

**StorageWorks**

Digital's family of modular data storage products that allows customers to design and configure their own storage subsystems. Components include power, packaging, cabling, devices, controllers, and software. Customers can integrate devices and array controllers in StorageWorks enclosures to form storage subsystems.



**StorageWorks building block**

See SBB.

**supported device**

A device tested as functionally compatible with an approved StorageWorks hardware and software configuration.

**warm swap**

A controller function that allows devices to be added, removed, or replaced while the subsystem remains operational. All activity on the device's SCSI bus must normally be halted for the duration of the warm swap operation.



## –A–

---

- Access to storage subsystem, 2-5
- Accessories
  - selecting, 4-16
- Adapters
  - selecting, 4-2
- Analyzing storage requirements, 2-1
- Application analysis
  - worksheet, 2-4
- Application requirements
  - access to storage subsystem, 2-5
  - data availability, 2-8
  - I/O, 2-7
  - storage capacity, 2-5
- Applications
  - access requirements, 2-5
  - data availability requirements, 2-8
  - I/O requirements, 2-7
  - I/O transfer requirements
    - data transfer-intensive, 2-7
    - I/O request-intensive, 2-7
    - read-intensive, 2-7
    - write-intensive, 2-7
  - storage capacity requirements, 2-5
- Availability, data
  - definition, 2-8
  - improvement of
    - using RAID, 3-11
    - using redundant controllers, 3-9
    - using redundant power systems, 3-9
  - planning for, 3-9
    - using redundant controllers, 3-9
    - using redundant power systems, 3-9

## –C–

---

- Cache
  - read, 3-5
  - write-back, 3-5
- Capacity planning, 3-4
- Configuration analysis
  - summary worksheet, example, 2-12
  - worksheet, example, 2-12
- Configuration Guide

- organization of, 1-4
- Configuration planning, 3-1
  - capacity planning, 3-4
  - for data availability, 3-9
    - cost/performance tradeoffs, 3-14
    - using RAID, 3-11
    - using redundant controllers, 3-9
    - using redundant power systems, 3-9
  - for subsystem performance, 3-5
    - cost/availability tradeoffs, 3-14
    - using controller caching, 3-5
    - using independent controllers, 3-7
    - using RAID, 3-11
  - summary worksheet, 2-15
- Configuration rules
  - for power systems, A-4
  - for SCSI bus, A-4
    - computing bus length, A-5
    - terminating, A-5
  - for shelves
    - general, A-1
    - with 3 port controllers
      - using 3.5-inch SBBs, A-2
      - using 5.25-inch SBBs, A-2
    - with 6 port controllers
      - using 3.5-inch SBBs, A-3
      - using 5.25-inch SBBs, A-3
- Controllers
  - cache
    - read, 3-5
    - write-back, 3-5
  - selecting, 4-2

## –D–

---

- Data availability
  - definition, 2-8
  - improvement of
    - using RAID, 3-11
    - using redundant controllers, 3-9
    - using redundant power systems, 3-9
  - planning for, 3-9
    - using redundant controllers, 3-9
    - using redundant power systems, 3-9

- Data categories
  - active, 2-6
  - dormant, 2-6
  - inactive, 2-6
  - recommended storage types
    - for nearline data, 2-6
    - for off-line data, 2-6
    - for on-line data, 2-6
- Device access time, 2-8
- E–**

---

- Enclosures
  - selecting, 4-15
- I–**

---

- I/O requirements, 2-7
  - device access time, 2-8
- Independent controllers, multiple, 3-7
- Introduction, 1-1
- M–**

---

- Magnetic disk, 4-5
- N–**

---

- Nearline storage, 2-6
- O–**

---

- Off-line storage, 2-6
- On-line storage, 2-6
- Optical disks, 4-5
- P–**

---

- Performance, subsystem
  - improvement of
    - using RAID, 3-11
  - improving
    - using controller caching, 3-5
    - using independent controllers, 3-7
  - planning, 3-5
    - for independent controllers, multiple, 3-7
    - using controller caching, 3-5
- Planning configurations, 3-1
- Power components
  - selecting, 4-15
- Power system configuration rules, A-4
- Power systems
  - redundant, 3-9
- R–**

---

- RAID, 3-11
  - advantages, B-1
  - level 0, B-2
  - level 0&1, B-6
  - level 1, B-3
  - level 3, B-4
  - level 5, B-5
  - number of member disks, 4-6
  - overview, B-1
  - sparesets, 3-11
  - summary comparison, B-7
- Redundant controllers, 3-9

- Redundant power systems, 3-9
- S–**

---

- SCSI
  - bus configuration rules, A-4
    - computing bus length, A-5
    - terminating SCSI buses, A-5
- Shelves
  - configuration rules
    - for 3 port controllers
      - using 3.5-inch SBBs, A-2
      - using 5.25-inch SBBs, A-2
    - for 6 port controllers
      - using 3.5-inch SBBs, A-3
      - using 5.25-inch SBBs, A-3
  - general, A-1
  - selecting, 4-8
    - for 3.5-inch devices, 4-9
    - for 5.25-inch devices, 4-10
- Storage capacity requirements, 2-5
- Storage capacity, definition, 3-4
- Storage devices
  - magnetic disk, 4-5
  - optical disk, 4-5
  - selecting, 4-5
  - solid state disk, 4-5
  - tape drives, 4-5
  - tape libraries, 4-5
  - tape loaders, 4-5
- Storage relationships, 3-4
- Storage requirements
  - analyzing, 2-1
- Storage types
  - nearline, 2-6
  - offline, 2-6
  - on-line, 2-6
- StorageWorks configuration
  - analyzing storage requirements, 2-1
  - implementing, 5-1
    - configuration rules
      - for power systems, A-4
      - for SCSI bus, A-4
        - computing bus length, A-5
        - terminating SCSI buses, A-5
    - for shelves
      - general, A-1
      - with 3 port controllers
        - using 3.5-inch SBBs, A-2
        - using 5.25-inch SBBs, A-2
      - with 6 port controllers
        - using 3.5-inch SBBs, A-3
        - using 5.25-inch SBBs, A-3
  - planning, 3-1
- StorageWorks products
  - selecting accessories, 4-16
  - selecting controllers and adapters, 4-2
  - selecting enclosures, 4-15

- selecting power components, 4-15
- selecting shelves for 3.5-inch devices, 4-9
- selecting shelves for 5.25-inch devices, 4-10
- selecting storage devices, 4-5
- Subsystem performance
  - improvement of
    - using RAID, 3-11
  - improving
    - using controller caching, 3-5
    - using independent controllers, 3-7
  - planning, 3-5
    - for independent controllers, multiple, 3-7
    - using controller caching, 3-5
- Summary worksheet, 2-15

**-T-**

---

Tape drives, 4-5

Tape libraries, 4-5

Tape loaders, 4-5

Terminating SCSI buses, A-5

Tradeoffs

- between cost/availability/performance, 3-14

**-U-**

---

User analysis

- worksheet, 2-3

**-W-**

---

Worksheets

- application analysis, 2-4
- configuration analysis, example, 2-12
- summary worksheet, 2-15
- summary, example, 2-12
- user analysis, 2-3