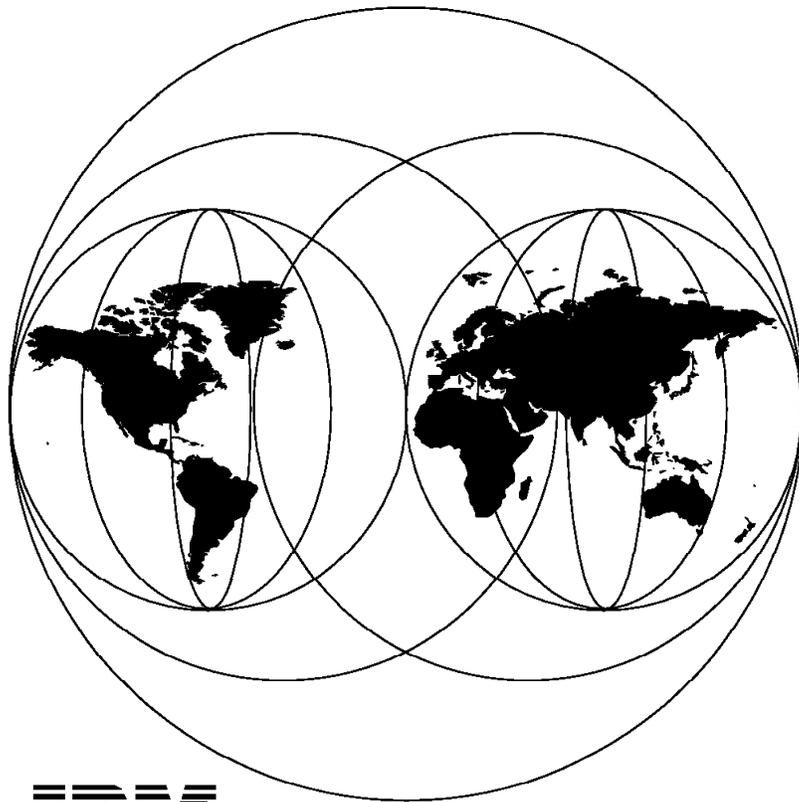


International Technical Support Organization

SG24-4599-00

**A Practical Guide to
Serial Storage Architecture for AIX**

July 1996



IBM

**International Technical Support Organization
Austin Center**



International Technical Support Organization

SG24-4599-00

**A Practical Guide to
Serial Storage Architecture for AIX**

July 1996

Take Note!

Before using this information and the product it supports, be sure to read the general information under "Special Notices" on page xiii.

First Edition (July 1996)

This edition applies to the IBM 7133 models 500 and 010, feature 6214 of the RS/6000 and AIX Version 3.2.5.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

An ITSO Technical Bulletin Evaluation Form for reader's feedback appears facing Chapter 1. If the form has been removed, comments may be addressed to:

IBM Corporation, International Technical Support Organization
Dept. JN9B Building 045 Internal Zip 2834
11400 Burnet Road
Austin, Texas 78758-3493

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1996. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Abstract

This document provides comprehensive information on Serial Storage Architecture for the RS/6000 with AIX Version 3.2.5. The technology is covered in some detail along with descriptions of the features and functions of the products currently available. Step-by-step instructions on the installation and configuration of the hardware and software are presented, followed by examples of various hardware and software scenarios tailored for common environments. Sections are also included on problem determination and recovery and performance and availability optimization.

This document is intended to assist systems engineers and customers to rapidly make the most effective use of Serial Storage Architecture on the RS/6000. The scenarios included can easily be implemented to immediately address most storage requirements, while the detailed technology information can be used to tailor configurations further in the future. Some knowledge of AIX Version 3.2.5 is assumed.

(170 pages)

Contents

Abstract	iii
Special Notices	xiii
Preface	xv
How This Document is Organized	xv
Related Publications	xvi
International Technical Support Organization Publications	xvi
How Customers Can Get Redbooks and Other ITSO Deliverables	xvi
How IBM Employees Can Get Redbooks and ITSO Deliverables	xviii
Acknowledgments	xviii
Chapter 1. SSA Technology Overview	1
1.1 History	1
1.1.1 Introduction	1
1.1.2 Connectivity	1
1.1.3 Serial Bus	2
1.1.4 IBM 9333 Serial Disks	2
1.2 Serial Storage Architecture	3
1.2.1 SSA Components	3
1.2.2 SSA Characteristics	4
1.3 The SSA Reality	11
Chapter 2. SSA Products Overview	13
2.1 The IBM 7133 Serial Storage Architecture Disk Subsystem	13
2.1.1 7133 Product Overview	13
2.1.2 7133 Technical Detail	14
2.2 Third-Party Serial Storage Architecture Products	31
Chapter 3. Installation and Configuration	33
3.1 General Procedure	33
3.1.1 Software Requirements	33
3.1.2 Hardware Requirements	35
3.2 Installing the SSA Software	37
3.2.1 Prerequisites	38
3.2.2 Actions	38
3.3 Configuration	39
3.3.1 Basic Configuration	39
3.4 Additional Configurations	42
3.4.1 Adding Disks	42
3.4.2 Removing Disks	44
3.4.3 Moving Disks	46
3.4.4 Changing the Path to a Disk	47
3.5 Administration of the subsystem	47
Chapter 4. Hardware Scenarios	49
4.1 Performance-Related Scenarios	49
4.1.1 Considerations	49
4.1.2 Scenarios	54
4.1.3 Summary	60
4.2 Availability Related Scenarios	60

4.2.1	Considerations	60
4.2.2	Scenarios	63
4.3	Performance and Availability	72
4.3.1	Considerations	72
4.3.2	Summary	75
Chapter 5. Problem Determination and Resolution		77
5.1	SSA Service Aids	77
5.1.1	Starting the SSA Service Aids	77
5.1.2	Using the Service Aids	79
5.1.3	Locating Devices without the Service Aids	92
5.2	Problem Determination and Resolution Examples	92
5.2.1	Service Aid Errors	92
5.2.2	Using the Service Aids to Resolve Problems	94
Chapter 6. Performance/Availability Investigation		99
6.1	Bandwidth	99
6.1.1	Individual Disk	100
6.1.2	SSA Loop	101
6.2	Data Path	102
6.3	Multiple Adapters	105
6.4	Data Location	106
6.5	Multiple Disk Types	106
6.6	Tuning SSA Parameters	107
Appendix A. SSA Product Reference Information		109
A.1	IBM 7133 SSA Disk Subsystem Model 010	109
A.1.1	Physical Specifications	109
A.1.2	Operating Environment	109
A.1.3	Limitations	110
A.1.4	Hardware Requirements	110
A.1.5	Software Requirements	110
A.1.6	Electronic Customer Support	111
A.1.7	Publications	111
A.1.8	Features	111
A.1.9	Feature Exchanges	138
A.1.10	Accessories	138
A.1.11	Customer Replacement Parts	138
A.1.12	Machine Elements	138
A.1.13	Supplies	138
A.1.14	Diskettes	139
A.2	IBM 7133 SSA Disk Subsystem Model 500	139
A.2.1	Physical Specification	139
A.2.2	Operating Environment	139
A.2.3	Limitations	140
A.2.4	Hardware Requirements	140
A.2.5	Software Requirements	140
A.2.6	Electronic Customer Support	140
A.2.7	Publications	140
A.2.8	Features	141
A.2.9	No Charge Specify Codes	141
A.2.10	Chargeable Feature Codes	142
A.2.11	Feature Exchanges	159
A.2.12	Accessories	159
A.2.13	Customer Replacement Parts	159

A.2.14 Machine Elements	159
A.2.15 Supplies	159
A.2.16 Diskettes	159
Appendix B. 7133 Configuration Planning Charts	161
B.1 Planning Chart	161
B.2 Configuration Planning Tables	162
B.2.1 Adapter Locations and Connections	162
B.2.2 7133 Subsystem Configuration and Connections	163
B.2.3 Topology Configuration Chart	164
List of Abbreviations	165
Index	167

Figures

1.	An SSA String	4
2.	An SSA loop	5
3.	An SSA Switched Network	5
4.	Frame format	8
5.	Flow Control	9
6.	An Example of Spatial Reuse	11
7.	IBM 7133 SSA Subsystems	13
8.	IBM 7133 Lights and Power	15
9.	IBM 7133 Component Location	16
10.	IBM 7133 Back Connectors	17
11.	IBM 7133 Lights and Power	18
12.	IBM 7133 Component Location	19
13.	IBM 7133 Back Connectors	21
14.	SSA Disk Drive Module	22
15.	SSA Four-Port Adapter	23
16.	IBM 7133 Hardware Schematic	24
17.	7133 Disk Numbering	27
18.	Dual 7133 Configuration	28
19.	Listing the SSA Product Software	34
20.	Listing the SSA adapter device information	34
21.	The 7133 Model 010	36
22.	The 7133 Model 500	37
23.	SMIT Software Install Screen	39
24.	A Basic 7133 Configuration	40
25.	SSA Service Aids Menu	41
26.	SSA Link Verification Screen	42
27.	8 Disk Configuration	43
28.	Eight Disk Configuration	44
29.	Setting Service Mode	45
30.	Leaving Set Service Mode	46
31.	Disk Ownership Example	52
32.	Multiple Adapters Example	53
33.	Sequential Access, Capacity 20 GB	55
34.	Sequential Access, Capacity 60 GB	56
35.	Sequential Access, Capacity 120 GB	57
36.	Random Access, Capacity 20 GB	58
37.	Random Access, Capacity 60 GB	59
38.	Random Access, Capacity 120 GB	59
39.	One Host, One Adapter, Capacity 20 GB	64
40.	One Host, One Adapter, Capacity 60 GB	65
41.	One Host, One Adapter, Capacity 120 GB	65
42.	One Host, Two Adapters, Capacity 20 GB	66
43.	One Host, Two Adapters, Capacity 20 GB	67
44.	One Host, Three Adapters, Capacity 20 GB	68
45.	One Host, Two Adapters, Capacity 60 GB	69
46.	One Host, Two Adapters, Capacity 60 GB	70
47.	One Host, Two Adapters, Capacity 120 GB	71
48.	Configuration Planning Chart Page ___	161
49.	Topology Planning Chart Page ___	164

Tables

1.	7133 Power Supply Configuration Rules	14
2.	7133 Models 500 and 010 Ordered Alone	25
3.	7133 Model 010 Ordered as Part of a System	26
4.	Maximum Throughput to a Single Disk	100
5.	Maximum Throughput to an SSA Loop	102
6.	Effect of Disk Position on Throughput (10 disks)	103
7.	Effect of Disk Position on Throughput (nine disks)	104
8.	Effect of Disk Position on Throughput (eight disks)	104
9.	Effect of Multiple Adapters on Throughput	105
10.	Effect of Data Location on Throughput	106
11.	Effect of Concurrency on Throughput	107
12.	Publications shipped with the 7133	111
13.	No charge feature code summary	111
14.	Chargeable feature code summary for the model 010	113
15.	Publications shipped with the 7133	141
16.	No charge feature code summary	141
17.	Chargeable feature code summary for the model 500	142
18.	Adapter Location and Connections	162
19.	7133 Subsystem Configuration and Connections	163

Special Notices

This publication is intended to help systems engineers and customers make the most effective use of the Serial Storage Architecture products for AIX. The information in this publication is not intended as the specification of any programming interfaces that are provided by the IBM 7133 models, feature 6214 of the RS/6000 or AIX Version 3.2.5. See the PUBLICATIONS section of the IBM Programming Announcement for the IBM 7133 models, feature 6214 of the RS/6000 and AIX Version 3.2.5 for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program or service is not intended to state or imply that only IBM's product, program or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY 10594 USA.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The information about non-IBM (VENDOR) products in this manual has been supplied by the vendor, and IBM assumes no responsibility for its accuracy or completeness. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any performance data contained in this document was determined in a controlled environment, and therefore, the results that may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

Reference to PTF numbers that have not been released through the normal distribution process does not imply general availability. The purpose of including these reference numbers is to alert IBM customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution process.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

AIX
Micro Channel

IBM
OS/2

POWERparallel
RISC System/6000
SP

POWERstation
RS/6000
SP2

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

PC Direct is a trademark of Ziff Communications Company and is used by IBM Corporation under license.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Limited.

Microsoft, Windows, and the Windows 95 logo are trademarks or registered trademarks of Microsoft Corporation.

Adaptec	Adaptec Incorporated
Canon	Canon Kabushiki Kaisha
Crosstalk	Digital Communications Associates, Inc.
Dell	Dell Corporation
Fujitsu	Fujitsu Limited
Future Domain	Future Domain Corporation
Micropolis	Micropolis Corporation
Mitsumi	Mitsumi Denki Kabushiki Kaisha
Molex	Molex Incorporated
Philips	Philips Consumer Electronics Company
Samsung	Samsung Electronics Company, Limited
SCSI	Security Control Systems, Inc.
Siemens	Siemens Company
Tandem	Tandem Computers Inc.
Western Digital	Western Digital Corporation

Other trademarks are trademarks of their respective companies.

Preface

This document is intended to assist customers and systems engineers to quickly make effective use of Serial Storage Architecture for the RS/6000. The technology is explained in detail to provide a basis for understanding the ways in which the hardware and software can be used and to provide a background for the subsequent chapters. The various products currently available are also described in detail to allow an informed selection of the relevant hardware to be made. A step-by-step guide to installing and configuring SSA hardware and software is included, followed by a set of practical hardware and software scenarios that together will allow rapid utilization of the benefits of SSA. The scenarios are designed to be applicable to a range of common situations and can be expanded on as required. Sections on problem determination, performance and availability optimization will help in creating and maintaining an optimum SSA environment.

How This Document is Organized

The document is organized as follows:

- Chapter 1, "SSA Technology Overview"

This chapter provides an overview of the current Serial Storage Architecture technology. This will include information on the various components and their interrelationships as well as technical details of their operation and the terminology used. This chapter can be skipped by those readers who already understand the Serial Storage Architecture technology.

- Chapter 2, "SSA Products Overview"

This chapter provides detailed information on the various Serial Storage Architecture products that are currently available for the RS/6000. Descriptions of the products and their capabilities will be included along with a section about some of the third-party products. This chapter can be skipped by those readers who are already familiar with the current range of IBM SSA products.

- Chapter 3, "Installation and Configuration"

This chapter provides detailed information on the installation and configuration of the various components that comprise SSA for the RS/6000. Product prerequisites are listed, and the process required to set up and configure the hardware and software is described step-by-step. This chapter can be skipped by those readers who have already installed and configured SSA for the RS/6000.

- Chapter 4, "Hardware Scenarios"

This chapter investigates some of the possible hardware configurations. Various useful configurations will be described as well as the environments in which they might best be used. The intent is to allow a practical subset of the large number of possible configurations to be defined, allowing rapid and effective use to be made of the SSA products. This chapter should be viewed by all readers who are implementing an SSA solution for the first time, and it provides useful reference material for all readers

- Chapter 5, "Problem Determination and Resolution"

This chapter looks at isolating and resolving problems when using SSA. The various tools and processes available for error detection and recovery will be investigated and some of the more common issues described. This chapter should be viewed by all readers initially as it contains useful information for avoiding problems as well as for solving them.

- Chapter 6, “Performance/Availability Investigation”

This chapter looks at techniques for configuring the SSA environment to maximize performance and availability. Optimum hardware configurations will be investigated, and the various software parameters that can be tuned will be discussed. This chapter should be viewed by those readers wishing to optimize the performance or availability of their SSA environment.

- Appendix A, “SSA Product Reference Information”

This appendix provides detailed, product-specific information on the SSA hardware. The various options and features available are listed for reference purposes. This appendix can be skipped by those readers not requiring detailed, product-specific information on hardware features.

- Appendix B, “7133 Configuration Planning Charts”

This appendix provides several planning charts that can be used to assist in the processes of designing and documenting the configuration of an SSA storage environment. Accurate documentation of an SSA configuration is essential to the effective management and administration of the subsystems.

Related Publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

- *7133 SSA Disk Subsystem: Operator Guide*, GA33-3259-00
- *7133 SSA Disk Subsystem: Service Guide*, SY33-0185-00
- *7133 Model 010 SSA Disk Subsystem: Installation Guide*, GA33-3260-00
- *7133 Model 500 SSA Disk Subsystem: Installation Guide*, GA33-3263-00

International Technical Support Organization Publications

- *AIX Storage Management*, GG24-4484-00

A complete list of International Technical Support Organization publications, known as redbooks, with a brief description of each, may be found in:

International Technical Support Organization Bibliography of Redbooks, GG24-3070.

How Customers Can Get Redbooks and Other ITSO Deliverables

Customers may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **IBMLINK**

Registered customers have access to PUBORDER to order hardcopy, to REDPRINT to obtain BookManager BOOKs

- **IBM Bookshop** — send orders to:

usib6fpl@ibmmail.com (USA)
bookshop@dk.ibm.com (Outside USA)

- **Telephone orders**

1-800-879-2755 (USA)	0256-478166 (UK)
354-9408 (Australia)	32-2-225-3738 (Belgium)
359-2-731076 (Bulgaria)	1-800-IBM-CALL (Canada)
42-2-67106-250 (Czech Republic)	45-934545 (Denmark)
593-2-5651-00 (Ecuador)	01805-5090 (Germany)
03-69-78901 (Israel)	0462-73-6669 (Japan)
905-627-1163 (Mexico)	31-20513-5100 (Netherlands)
064-4-57659-36 (New Zealand)	507-639977 (Panama)
027-011-320-9299 (South Africa)	

- **Mail Orders** — send orders to:

IBM Publications
P.O. Box 9046
Boulder, CO 80301-9191
USA

IBM Direct Services
Sortemosevej 21,
3450 Allerød
Denmark

- **Fax** — send orders to:

1-800-445-9269 (USA)	0256-843173 (UK)
32-2-225-3478 (Belgium)	359-2-730235 (Bulgaria)
905-316-7210 (Canada)	42-2-67106-402 (Czech Republic)
593-2-5651-45 (Ecuador)	07032-15-3300 (Germany)
03-69-59985 (Israel)	0462-73-7313 (Japan)
31-20513-3296 (Netherlands)	064-4-57659-16 (New Zealand)
507-693604 (Panama)	027-011-320-9113 (South Africa)

- **1-800-IBM-4FAX (USA only)** — ask for:

Index # 4421 Abstracts of new redbooks
Index # 4422 IBM redbooks
Index # 4420 Redbooks for last six months

- **Direct Services**

Send note to softwareshop@vnet.ibm.com

- **Redbooks Home Page on the World Wide Web**

<http://www.redbooks.ibm.com/redbooks>

- **E-mail (Internet)**

Send note to redbook@vnet.ibm.com

- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

How IBM Employees Can Get Redbooks and ITSO Deliverables

Employees may request ITSO deliverables (redbooks, BookManager BOOKs, and CD-ROMs) and information about redbooks, workshops, and residencies in the following ways:

- **PUBORDER** — to order hardcopies in USA

- **GOPHER link to the Internet**

Type GOPHER

Select IBM GOPHER SERVERS

Select ITSO GOPHER SERVER for Redbooks

- **Tools disks**

To get LIST3820s of redbooks, type one of the following commands:

```
TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET GG24xxxx PACKAGE
```

```
TOOLS SENDTO CANVM2 TOOLS REDPRINT GET GG24xxxx PACKAGE (Canadian users only)
```

To get lists of redbooks:

```
TOOLS SENDTO WTSCPOK TOOLS REDBOOKS GET REDBOOKS CATALOG
```

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET ITSOCAT TXT
```

```
TOOLS SENDTO USDIST MKTTOOLS MKTTOOLS GET LISTSERV PACKAGE
```

To register for information on workshops, residencies, and redbooks:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1996
```

For a list of product area specialists in the ITSO:

```
TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ORGCARD PACKAGE
```

- **Redbooks Home Page on the World Wide Web**

<http://w3.itso.ibm.com/redbooks/redbooks.html>

- **ITSO4USA category on INEWS**

- **IBM Bookshop** — send orders to:

USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

- **Internet Listserver**

With an Internet E-mail address, anyone can subscribe to an IBM Announcement Listserver. To initiate the service, send an E-mail note to announce@webster.ibm.com with the keyword subscribe in the body of the note (leave the subject line blank). A category form and detailed instructions will be sent to you.

Acknowledgments

This project was designed and managed by:

Nick Higham

International Technical Support Organization, Austin Center

The authors of this document are:

Nick Higham

IBM Austin

Tony Sealy
IBM UK

This publication is the result of a residency conducted at the International Technical Support Organization, Austin Center.

Thanks to the following people for the invaluable advice and guidance provided in the production of this document:

Al Waisanen
SSD, IBM Austin

Steve Furniss
IBM UK

Additional thanks to

Marcus Brewer
International Technical Support Organization, Austin Center

Chapter 1. SSA Technology Overview

In this chapter, the basic elements of Serial Storage Architecture (SSA) are discussed. This will include the history behind the development of SSA, its components and characteristics and associated terminology.

1.1 History

This first section looks at the history behind the evolution of the Serial Storage Architecture, thereby putting into context the design decisions.

1.1.1 Introduction

In the storage subsystems arena, it is recognized that the Small Computer System Interface (SCSI) as a fast low-cost external interface attachment has a number of limitations:

- Speed improvement would involve increasing the data width, which would basically mean doubling it; this would result in:
 1. Cable and connector doubling in size
 - As the cables are usually low cost (ribbon or twisted pair), they tend to be bulky by default.
 2. Crosstalk
 3. An increase in the risk of disturbance caused in one data signal by an unwanted transfer of energy from another data signal
- Electromagnetic interference (EMI) and data deviation complications are limitations to using cheap unshielded cables.

The SCSI specification has introduced versions (such as SCSI-2) which increase its clock speed and data width, allowing it to operate at 20 MB per second, but it is recognized that further expansion possibilities are limited.

1.1.2 Connectivity

An alternative approach to increasing the data rate by increasing the physical size of the cable was to *reduce* the size of the cable to a single signal. Having narrow communication paths usually involves converting the wider data units into sequential bit processing; these are referred to as *serial* links or buses.

This would seem to be a backward step given the performance advantages of a wide data bus. However, a serial I/O connection can have several advantages:

- The connecting cable has less wire. There are obvious cost and bulkiness advantages; it also allows the opportunity for sophisticated shielding to be applied to the cable. This would be more expensive on multi-signal cables. Also, many of these changing at the same time, as is typical of wide data buses, can cause severe *ground bounce* problems.
- A single data signal means that there is no clock deviation to consider as the data rate or the length of the cable increases.
- Other transmission mediums, like optical, radio and telephone, can be used to continue the I/O connection; these are not easy to use in parallel transfers.

- The Application Specific Integrated Circuits (ASICs) that implement the connection can have very few pins devoted to the connection itself. In most technologies, it is the number of pins rather than the number of logic cells that dictate the cost of the ASIC.
- The ASICs have less of their circuitry devoted to the drivers and receivers of the signals. These components are often in short supply, difficult to place on an LSI chip for optimum performance and usually consume a lot of power.

1.1.3 Serial Bus

Although in theory the idea of a serial link conjures up an image of one wire, the nature of the transmission medium, the speed of the data and the requirements of the connection protocol make this type of link a rarity; typically, the RS232 serial link has nine connections. There are some good reasons for the extra connections:

- A reliable transmission protocol usually demands two-way, or *full-duplex*, communication. So the second communication path can be used for acknowledgement, bidirectional concurrent data transfer or both. When this is used in a *loop* (see 1.2.2.3, “Loops” on page 5), the second path may be used to double the bandwidth to the device or provide an alternative route to the device should one path fail.
- The most effective transmission of high frequency signals using copper wire is achieved by using *differential pairs* where the same signal is transmitted with the opposite polarity on the other wire of the pair. The two wires experience similar noise and distortion corruption during transmission, so the signal may be extracted by subtracting the output of one member of the pair from the other.
- Some serial links transmit the clock on a second signal. This means that the reception of the data is simplified, and variation of the clock speed may be accomplished transparently (with clock deviation considerations). However, the addition of a second signal for a clock represents a doubling of the signals required.

Examples of serial buses:

- Philips Inter-Integrated-Circuit bus
This is a 100 Kb per second, two-signal (clock and data) bus for connecting integrated circuits.
- IEEE P1394
This is a 393 Mb per second, two-signal (data and strobe) point-to-point connection for desktop computers.
- Fiber Channel (FCS)
This is a 1 Gb per second high-end optical link for local area networks.

1.1.4 IBM 9333 Serial Disks

In July 1991, IBM announced the 9333 High-Performance Disk Drive Subsystem, an external disk subsystem using a serial link architecture with the following features:

- Very high sustained I/O rate (1000 disk operations per second)
- Exploitation of 40/80 MB per second streaming on the RS/6000 Micro Channel

- Mixing of data and commands in full duplex stream allowing multiple concurrent operations
- Elimination of bus contention and maintenance of high performance in high capacity systems
- Expansion without performance degradation

The 9333 was the first generation of serial disks, but it was never designed as an industry standard, only as a proprietary subsystem for RISC/6000 and AIX.

1.2 Serial Storage Architecture

Serial Storage Architecture (SSA) is a serial link designed especially for low-cost high-performance connection to disk drives. It is a two-signal connection (transmit and receive) providing full duplex communication. It uses a self-clocking code, which means that the data clock is recovered from the data signal itself rather than being transmitted as a separate signal. For transmission along a copper wire, it uses the differential pair method, requiring four wires, but it can also be transmitted along fiber optic cable.

The copper wire transmission rate is currently 20 MB per second in each direction, for a maximum of 25 meters cable length.

1.2.1 SSA Components

Serial Storage Architecture currently comprises three components:

SSA-PH1	The physical layer. This is the physical and electrical specification of the serial link, drivers and connectors. of this chapter.
SSA-TL1	The transport layer. This layer defines the messages and control information for transmission of data across SSA. of this chapter.
SSA-S2P	The SCSI-2 mapping. It is recognized that the <i>logical</i> aspects of the SCSI specification remain appropriate for addressing serially attached peripherals and so this aspect of SCSI has been mapped to the physical SSA interface. This means that a transition may be achieved from SCSI to SSA with the absolute minimum of code rewriting.

In addition, it is planned that similar mappings will be provided for the SCSI-3 specification when it is approved by ANSI, and for higher data rates. The approval of any standard by ANSI is a long process and at the time of writing, this standard has not yet completed the public review part of the process. All SSA products now available from IBM comply with the 1995 level of the architecture as defined in the following documents:

- SSA-IA95/PH (defines SSA-PH1 and SSA-TL1)
- SSA-IA95/SP (defines SSA-S2P)

1.2.2 SSA Characteristics

The SSA serial link has the following characteristics:

- Topology
A flexible addressing scheme that allows connections to be made as strings, loops and switched loops. *Hot plugging* is also supported.
- Distance
Point-to-point connection for up to 25 meters (cable). Fiber-optic connections could support distances of up to 2.4 kilometers between nodes.
- Bandwidth
Full duplex communication, 20 MB per second in each direction. The protocol allows for speed increases as technology permits, in powers of two.
- Format
The unit of transmission is the frame, which can be up to 128 bytes. The minimum overhead per frame is eight bytes, or about six percent.
- Reliability
The link is highly reliable. The design provides for accurate detection of errors of all kinds and a large amount of transparent error recovery.
- Physical
The cables and connectors have a small form factor.

1.2.2.1 Topology

The SSA design allows an extremely flexible assortment of connection options. SSA networks can be connected in simple strings, loops, more complex switched strings, or cyclic paths. This flexibility allows trade-offs to be made between cost, performance and availability.

This variation is afforded by three different types of SSA node:

- Single port
- Dual port
- Switch

1.2.2.2 Strings

A *string* is a simple linear network of two or more nodes, as shown in Figure 1. The nodes at either end can be single-port nodes, while the others must be dual-port nodes.

A special case of a string is the *dedicated connection*, where two single-port nodes connect to each other across one link.

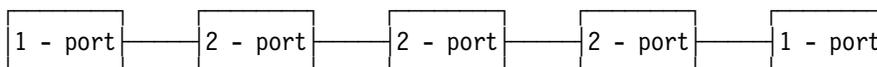


Figure 1. An SSA String

1.2.2.3 Loops

The most common form of connection is the *loop*, shown in Figure 2. A loop is a cyclic network containing only dual-port nodes. Loops have the benefit of both higher bandwidth and higher reliability over strings:

- Higher bandwidth
There are two data paths available between any two nodes.
- Higher reliability
Any single node can fail without prohibiting communication between nodes, and a node may be inserted into the loop without breaking communication.

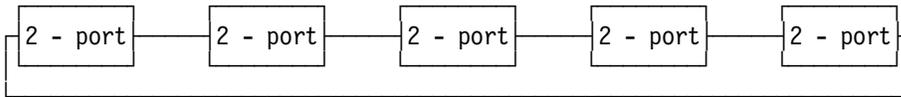


Figure 2. An SSA loop

1.2.2.4 Switches

Figure 3 shows an example of a complex network involving a *switch*.

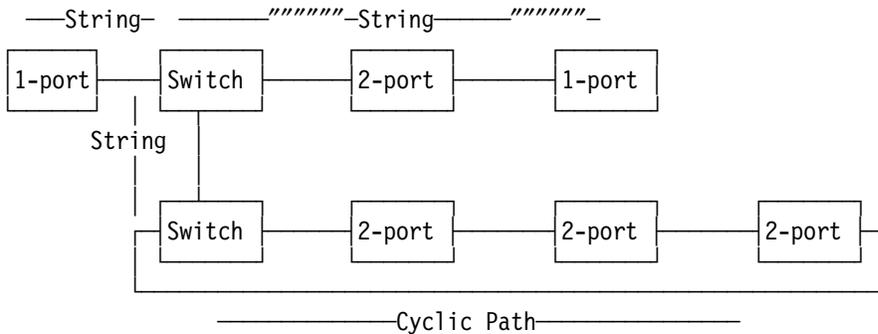


Figure 3. An SSA Switched Network

A switch can have up to 126 ports. Switches allow large numbers of nodes to be connected together and also enable alternative paths to be established to provide fault tolerance.

A switched network can also include other *cyclic paths*; these are not loops, by definition, because they involve types of nodes other than dual-port nodes.

1.2.2.5 Coding

Despite the serial nature of the link, the unit of information transmitted remains the byte. Conventionally, an 8-bit byte is *serialized* and sent as a bit stream. However, if these bytes were serialized and sent one after another, with no intervening control bits (as would be necessary for maximum throughput), the following problems would occur:

1. There would be no way to recover and synchronize the data clock, especially during a long stream of zeros, for example.
2. Any errors occurring would be undetected and uncorrected.
3. Whenever more bits of one polarity than the other are sent, there will be a DC bias in the data signal.

To overcome these problems, the serialized bit stream is usually at a fixed clock rate, and interspersed with control bits (such as stop, start and parity), which helps to resolve the first three problems at the expense of some bandwidth.

The DC bias is usually ignored, even though it is sometimes an inhibitor to increasing speed over the link. This is because as the data rate increases, the data signal contains more components at higher frequencies. To extract these signals reliably sometimes requires AC-coupled amplifier circuits at the receiver end.

An example of this would be where the average value of a DC free signal (obtained by simple integration) can be used to provide the *slicing* point, the level that distinguishes a one from a zero. This is particularly useful for fiber-optic receivers and is also the principle used in coding on compact disks and CD-ROMs.

Due to these considerations, SSA uses a form of coding called *8B/10B* encoding. The name reflects the fact that the eight bits to be coded become 10 bits of data. There is a 25 percent redundancy in the signal, but this overhead is well spent.

A 10-bit value can take on up to 1024 different combinations, but some of these may not be valid due to the following rules:

- There can be no more than five consecutive bits of the same value (this applies across adjacent bytes as well as within a byte).
- The maximum *Digital Sum Variation (DSV)* is six (+3 through to -3).

The run-length limit ensures that the clock can be successfully recovered at all times and will remain synchronized.

The DSV is defined as follows: counting 1 as positive, and zero as a negative, a running count is kept as the bits are encoded. Taking the maximum value for the count and subtracting the minimum value produces the DSV. If this value is constrained, then the DC component of the transmitted signal is effectively zero because there are as many positives as there are negatives in the overall stream.

To achieve this, the codes selected for transmission are dependant on the codes that have previously been used as well as the data that is to be transmitted. Thus, each byte to be encoded can map to more than one code that is transmitted; this accounts for the redundancy in the coding method.

In practice, the 8-bit byte to be encoded maps to one or two codes. The byte is split into two parts, a 5-bit string and a 3-bit string. The substrings are then encoded using a 5B/6B code for the first substring and a 3B/4B code for the second. Each substring has either an exact number of positives and negatives or is unbalanced by two. An imbalance is corrected by selecting the version of the next byte's code that has the opposite imbalance.

This coding method successfully encodes all 256 possible values for a byte and allows a further 12 codes to be included that obey the coding rules. Of these, three have the unique property of containing a string of bits within them (00111111) that can never occur in any other bit position, either within a code or across a boundary between two words. These are known as *comma* codes, and they may be used by the receiver to synchronize its clock to the correct

codeword boundary. The remaining codes are used by the SSA protocol as special characters because they are readily identified as non-data.

The special characters in SSA are:

FLAG	Used to delimit the start and end of a <i>frame</i> (see 1.2.2.6, “Frames”). It is also a comma character.
ABORT	Prematurely aborts the transmission of a frame.
SAT	Used for arbitration With <i>SAT</i> . See 1.2.2.11, “SAT Algorithm” on page 11.
SAT’	Used for arbitration with <i>SAT</i> . See 1.2.2.11, “SAT Algorithm” on page 11.
DIS	Indicates the disabled state. It is also a comma character.
ACK	Acknowledges correct reception of a frame.
RR	Indicates readiness to accept another frame.
NUL	Used to pad out transmission, for example when waiting for some response before proceeding. This character is ignored.

There is one other comma character which, under certain circumstances, can cause erroneous detection of a comma, so this character is declared invalid.

The three remaining special characters are available for use by the application or the higher level protocol. One of these has been defined for SSA-S2P as a *SYNC* character, used to synchronize the rotation of the disk drives. The other two are currently reserved.

The advantages of the 8B/10B code are:

- The clock is always recoverable. The maximum number of bits without a transition is five.
- No DC content, allowing AC-coupled signal processing.
- Some inherent error checking. With so many of the combinations and sequences being invalid, many of the errors due to noise may be detected just by verifying the coding rules.
- Special characters are available; non-data characters that can be used for special signalling (even within data sequences) and byte synchronization.

1.2.2.6 Frames

Data for transmission is organized into *Frames*. A frame consists of a sequence of at least six data characters delimited by a *FLAG* at each end. The components of a frame are:

Control	Single data character. This is used to indicate the type of frame, which can be: <ul style="list-style-type: none">• Control A control frame issues various types of reset.• Privileged A privileged frame is used for certain protocol related messages.
----------------	--

- Application

An application frame is used to send normal data.

The control byte is also used to indicate the *frame sequence number*, this is used to detect lost frames.

- Address** Consists of one to six data characters. The address field is used to direct a frame to the correct node in a string, loop or switch, using a *path* value and also to direct the frame to the correct *channel* within that node.
- Data** Consists of zero to 128 characters. This field contains the application data if the frame is an application frame or a privileged message for a privilege frame. It is not present for control frames. The length of the data field must be a multiple of 8 bytes, except for message frames and the frame containing the last byte of data. Message frames are a maximum of 32 bytes.
- CRC** Consists of four data bytes. The CRC field immediately precedes the trailing FLAG. It is calculated based on the control, address and data fields. It cannot, and does not, include the FLAG characters or any intervening special characters. It uses the degree 32 polynomial used by FDDI and FCS.

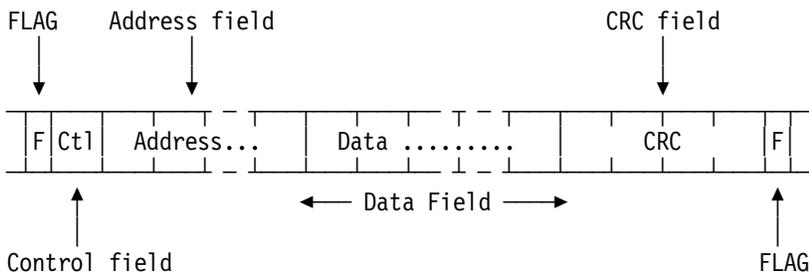


Figure 4. Frame format

1.2.2.7 Flow Control

SSA frames can be sent from any node to any other node in the network. On its journey, a frame may traverse several links. The integrity of the data across each link is the responsibility of the two nodes involved, even though they may only be couriers.

For the links to be reliable, there are tight controls on the *flow* of data between nodes. Each node must contain enough buffer space to hold at least one complete frame. After transmission, the transmitter must wait until it receives an *ACK* from the receiving node before it may discard the data. The *ACK* is only sent if the received frame passes all the error checks. The *ACK* signal is sent twice to ensure safe reception.

The error checks cannot be performed instantaneously, so in order to allow maximum bandwidth utilization, the receiver can (assuming there is sufficient buffer space) request a second frame before it has acknowledged the first frame. This is achieved by sending a *Receiver Ready (RR)* character back to the transmitter (see 1.2.2.5, "Coding" on page 5). This may be done as soon as

there is buffer space, so it may be sent before the transmitter has finished sending the first frame. The RR character is also sent twice in succession.

The RR character informs the transmitter that the second frame may be sent before the first frame is acknowledged; however, it still requires the ACK signal before it can discard its copy of the first frame. The ACK for the first frame must be received before a third frame can be sent, no matter how much buffer space the receiver has; this is so that the ACK may be correctly associated to the frame to which it applies. If the transmitter does not receive the ACK before it has finished transmitting the second frame, it must insert NUL characters before the trailing FLAG. Figure 5 shows the flow of ACKs and RRs.

The ACK and RR signals are only between individual nodes of a link; even if a node is simply passing data through to another link, it must obey the correct ACK and RR protocol. ACKs and RRs are not forwarded, they are generated by each receiver on an individual basis.

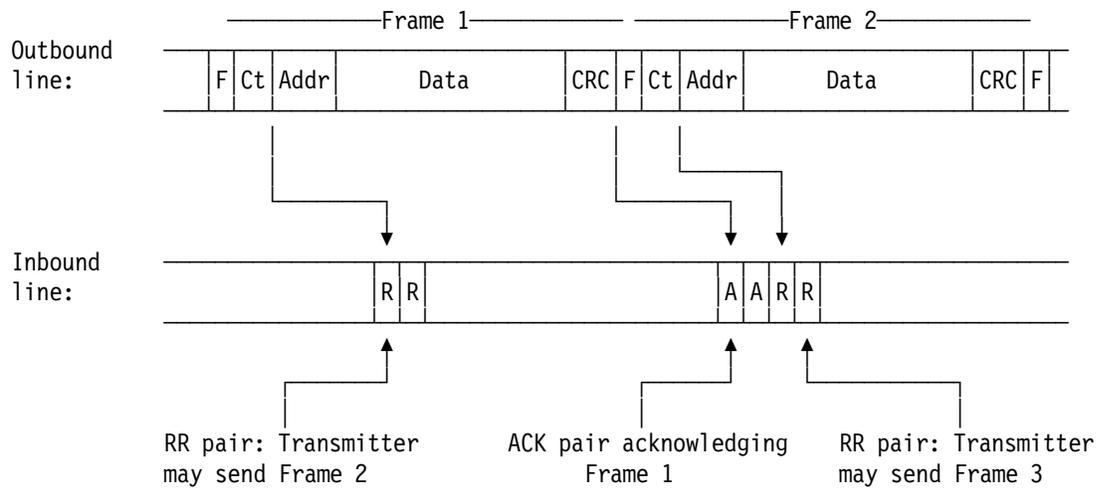


Figure 5. Flow Control. The operation for half duplex is shown. For full duplex, the ACK and RR pairs would be interspersed within the frames on the inbound signal.

1.2.2.8 Cut-Through (Worm-Hole) Routing

It would appear at a glance that with all the error checks and flow control between links that a frame that crossed several links would make a very sluggish journey towards its intended destination, held up at each node while its credentials were checked. However, the expected error rate across SSA links is so low that such delays can be avoided by a method known as *cut-through routing*.

Cut-through routing means that a node may forward a frame character-by-character as it is received; it does not have to wait for confirmation that the frame passed its CRC check. Using this method, the delay can be as little as 5-10 characters or 0.5 microseconds at 20 MB per second.

If an error is detected on an inbound frame after a router has already started to forward the frame, then it sends an ABORT character (followed by a FLAG) to the receiver. The receiver will send an ABORT if it, too, has already begun forwarding. An ABORT character tells a receiver to discard the frame in which it occurs; however, the receiver that detected the error must now attempt to recover the error, as described in the next section.

1.2.2.9 Error Handling

Although the low signal count of a serial link affords good quality cable shielding, errors can occur and, if left uncorrected, their consequences, particularly where disk drives are concerned, can be disastrous.

In addition, the SSA architecture allows *hot plugging*, which means that nodes can be connected and disconnected without powering down or giving the other nodes any special warning. These *deliberate* errors must also be detected and corrected.

A large part of the reliability of SSA is due to its *Error Recovery Procedures (ERP)*. The features of the SSA design (8B/10B 32-bit code) make it particularly good at detecting errors due to noise; it is estimated that the undetected error rate is less than 1 in 1×10 to the power of 22 characters. An error can be detected by any of the following:

- ACK Time-out (50 microseconds)
- Loss of synchronization
- 8B/10B code violation
- Protocol error (for example an isolated single ACK)
- CRC error
- Frame sequence error
- Rejected frame (too many data characters, invalid address, etc)

as well as any internal errors caused by the links themselves.

Any of these errors occurring invokes the *Link ERP* which attempts to recover the error. If successful, the recovery is transparent to the higher level protocol, with a slight loss of bandwidth.

The link ERP will first reset the link, which causes the nodes concerned to exchange their status. This gives information about the type of error that occurred. If it is recoverable, the transmitter will *retransmit* the corrupted frame, as it will still have a copy (the ACK would only have been sent if the frame had passed all the checks). If this re-transmission is successful (before a defined number of retries), then the link continues with transmission without informing the higher level protocol; the error has been transparently recovered.

If the error was not recoverable, either because re-transmission failed or the error was too severe, the ERP *exits* with an *Asynchronous Alert*, which is one of the *privileged messages*. This marks the link as non-operational; however, in the case of a loop connection, traffic may be rerouted to avoid this link, one of the advantages of a full duplex link.

The correct protocol must be observed along *every* link. If a receiver detects an error in transmission it must attempt to correct that error, even though it is not the intended destination for a transmission (a courier for example). So each link will only forward good data and the failing link can always be identified.

1.2.2.10 Spatial Reuse

One of the characteristics that distinguishes a *ring* (where each link between nodes is a separate connection) and a bus (where each node connects to the same pieces of wire) is that a ring allows the possibility of *Spatial Reuse*. This is the technique whereby links that are not involved in a particular data transaction are available for use in another transaction. This is not possible on a bus because the entire bus carries the same information.

An example of spatial reuse is shown in Figure 6. Node D is transmitting to node C at the same time as node E is transmitting to node B. The concurrent transaction is possible because the links are independent.

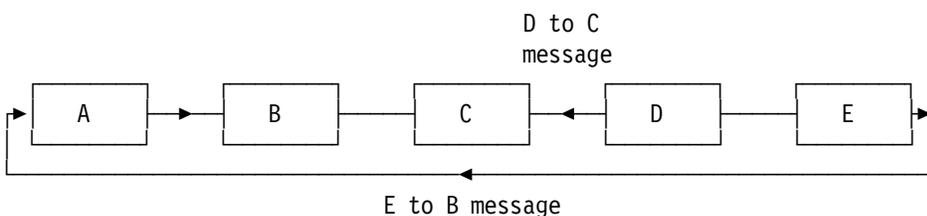


Figure 6. An Example of Spatial Reuse

1.2.2.11 SAT Algorithm

Fair spatial reuse is made possible by the use of an algorithm designed to provide a fair balance between the duties of originating frames and simply passing on frames meant for other nodes. In this technique, a token circulates in the *opposite* direction to the data transfer. This token, known as *SAT*, allows the possessor to originate an assigned minimum number of transmissions (*SAT*isfy itself) in that direction, after which it must pass the token on. When not in possession, a node may originate up to its maximum number of transmissions, giving priority to frames it is passing on from other ports. The minimum and maximum numbers, known as the *A_Quota* and *B_Quota*, are programmed during configuration time.

This algorithm requires a bidirectional link, such as SSA, in order to operate. But to allow the fairness algorithm to be applied in both directions, SSA defines a complementary token *SAT'* to govern traffic in the opposite direction. *SAT* and *SAT'* are two of the special characters allowed by the 8B/10B coding and so are easily distinguished from other data characters. Consequently, they may be inserted almost anywhere in a data stream. A string network implements the algorithm by reflecting the token in its complementary form (*SAT* becomes *SAT'* and *SAT'* becomes *SAT*) in the nodes at the ends of the string.

Using this algorithm, an SSA loop has almost three times the throughput of a dual, counter-rotating token ring.

1.3 The SSA Reality

So far, only the concepts of SSA have been discussed. In the next chapter, the implementation of SSA design into specific products will be illustrated.

Chapter 2. SSA Products Overview

This chapter describes the IBM Serial Storage Architecture products available, which includes both rack-mounted and desktside models of the IBM 7133 SSA Subsystem and the associated host adapters, providing detailed information on their capabilities. Brief information is also provided on some third party products to allow a comparison to be drawn between available SSA products.

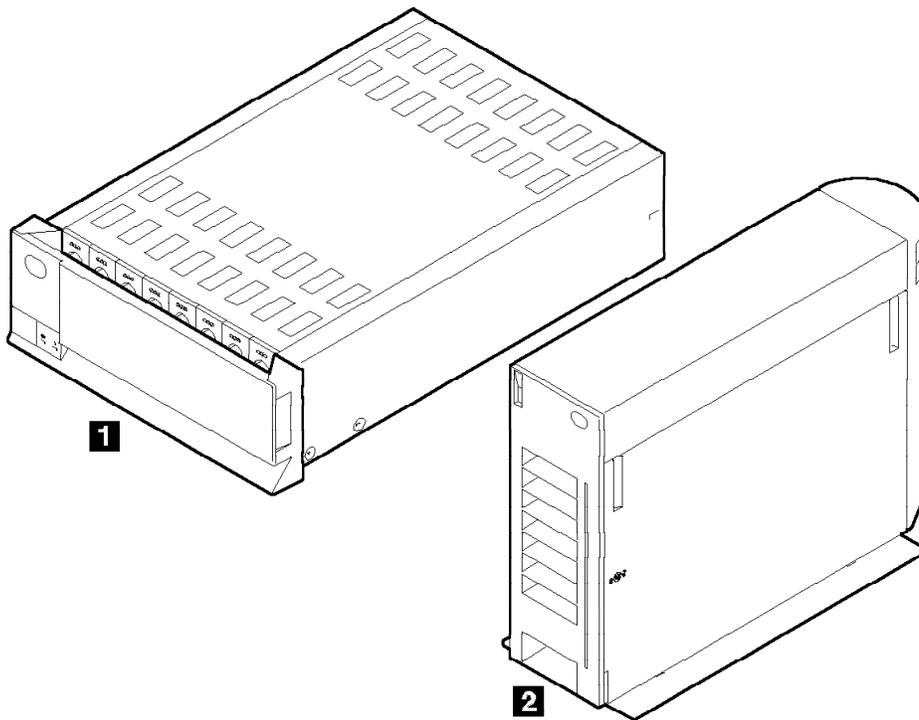


Figure 7. IBM 7133 SSA Subsystems. The rack-mounted and tower models of the IBM 7133 SSA subsystem.

2.1 The IBM 7133 Serial Storage Architecture Disk Subsystem

The IBM 7133 SSA Disk Subsystem provides an industry standard implementation of the Serial Storage Architecture for both the RS/6000 and RS/6000 SP hardware platforms. The 7133 is available in the two models depicted in Figure 7:

1. Rack-mounted model 010
2. Stand-alone tower model 500

2.1.1 7133 Product Overview

Both models are functionally equivalent, with the model 010 being designed to be integrated into an industry standard 19-inch rack, such as is used by the RS/6000 900 and R series, the RS/6000 SP models and the 7202 model 900 expansion rack. Three different capacity disk drives are available for use in both models of the 7133. These are the 1.1 GB, 2.2 GB and 4.5 GB Ultrastar drives. The disk drives are supplied mounted on an auto-docking carrier for easy

insertion and removal, and both models of the 7133 will support from four to 16 of them, providing a maximum capacity of 72 GB. In addition, the 7133 models include two auto-docking power and cooling units that provide redundant power and cooling for the first four disk drives and non-redundant power and cooling to the second set of four drives. The addition of the third power and cooling module ensures redundant power and cooling to all 16 drives.

Attachment of the 7133 subsystems to the RS/6000 host is achieved through the IBM SSA 4-port adapter (feature code 6214, or the newly announced SSA adapter, feature code 6216) for the RS/6000 and RS/6000 SP models.

The 7133 SSA Disk Subsystem is currently supported by AIX Version 3.2.5 with additional PTFs and with the AIX 3.2.5 device driver that is shipped with the SSA four port adapter (feature codes 6214 or 6216 on the RS/6000). The required PTFs and specific software prerequisites are listed in Chapter 3, "Installation and Configuration" on page 33. For convenience, the PTFs are packaged with the device driver.

2.1.2 7133 Technical Detail

This section looks first at the technical details of the 7133 Serial Storage Architecture Disk Subsystem and then at the rules governing its configuration. The third section looks in more detail at the software support for the subsystem and adapter.

2.1.2.1 7133 Physical Layout

The first part of this section looks at the physical layouts and component location of the two models of the 7133. This is intended to assist in insertion and removal of disks, power supplies and cables. The second part then shows the various components and describes the rules governing their interconnection.

There are certain rules governing the way in which power and cooling modules must be installed with disk modules, and these are summarized in Table 1.

<i>Table 1. 7133 Power Supply Configuration Rules</i>			
Configuration	Number Of Disk Modules	Minimum Requirements without Redundancy	Minimum Requirements with Redundancy
Base	4	Always redundancy	Power and cooling modules in slots 3 and 2
Expanded	5-8	Power and cooling modules in slots 3 and 2	Power and cooling modules in all slots
Maximum	9-16	Always redundancy	Power and cooling modules in all slots
<p>Note:</p> <ol style="list-style-type: none"> 1. Errors will occur if a dummy power supply and cooling module is not present in an empty power supply slot. 2. There will always be power and cooling modules in slots 3 and 2. 3. All empty drive module slots must contain dummy modules. 			

Physical layout of the Model 010: This section will illustrate the location of the various components and indicator lights in the 7133 model 010.

Lights and Power: The diagram in Figure 8 shows the location of the indicator lights and power switch in the rack-mounted 7133 model 010.

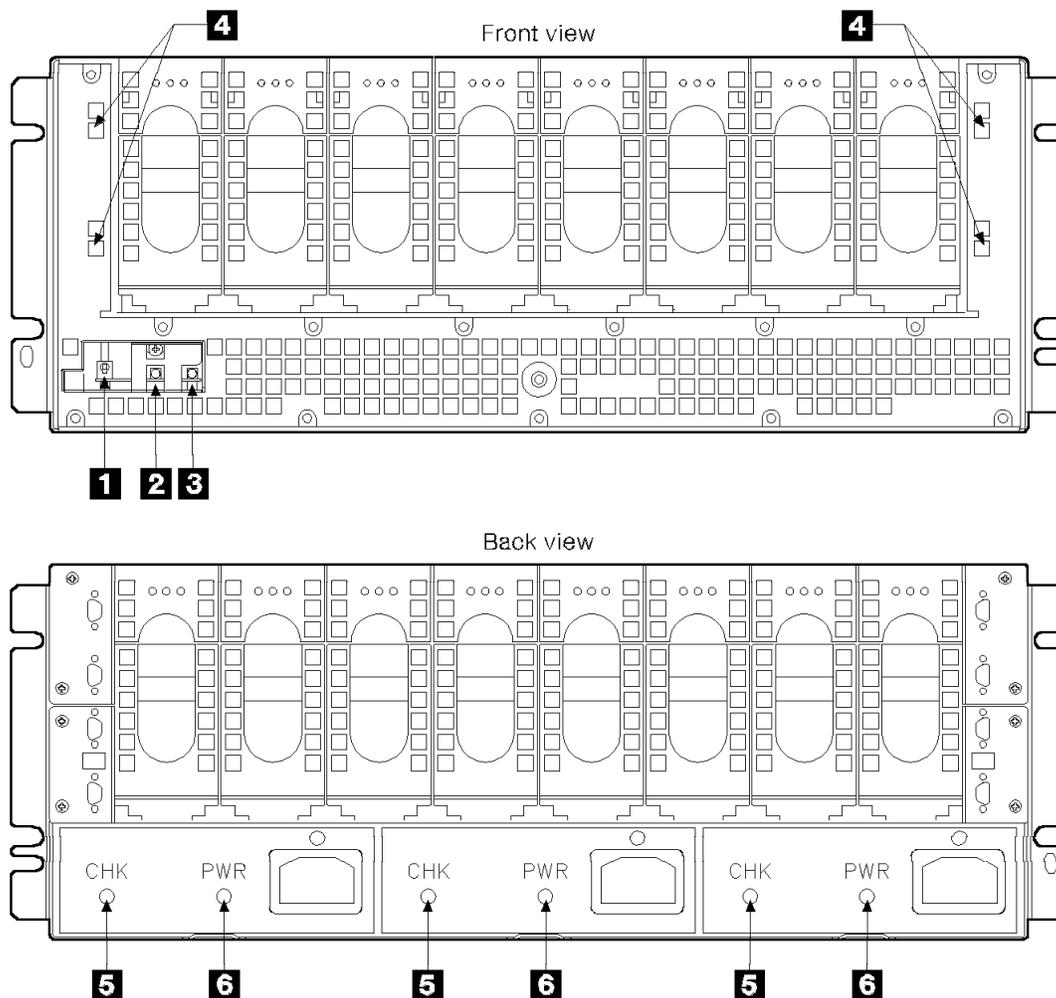


Figure 8. IBM 7133 Lights and Power. The location of the indicator lights and power switch.

1. Power switch

This switch controls the flow of internal DC power from the power and cooling modules to the rest of the unit. Power will always be present in the power and cooling modules if they are connected to an external power source. To activate the 7133, press and hold the switch for two seconds. To deactivate the 7133, again press and hold the switch for two seconds.

A flap over the switch prevents accidental operation. This flap pivots upwards to allow access to the switch.

2. Subsystem power light

When lit, this green light indicates that the 7133 is being supplied with DC power.

3. Subsystem check light

When lit, this amber light indicates that a failure has been detected within the 7133. The unit may still be able to operate normally, depending upon the nature of the problem.

4. Power card light

When lit, these green lights indicate that DC power is being supplied to the power cards for the disks.

5. Fan-and-power check light

When lit, these amber lights indicate that there is a problem with the DC power supply. If the remote-power-on feature is being used, this light indicates that power has been terminated by the using system or by the remote-power-on control unit.

If the light is flashing it indicates a cooling fan failure.

6. Power light

When lit, these green lights indicate that external AC power is being supplied to the power and cooling modules.

Parts Location: The diagram in Figure 9 shows the location and organization of the various components inside the rack-mounted 7133 model 010.

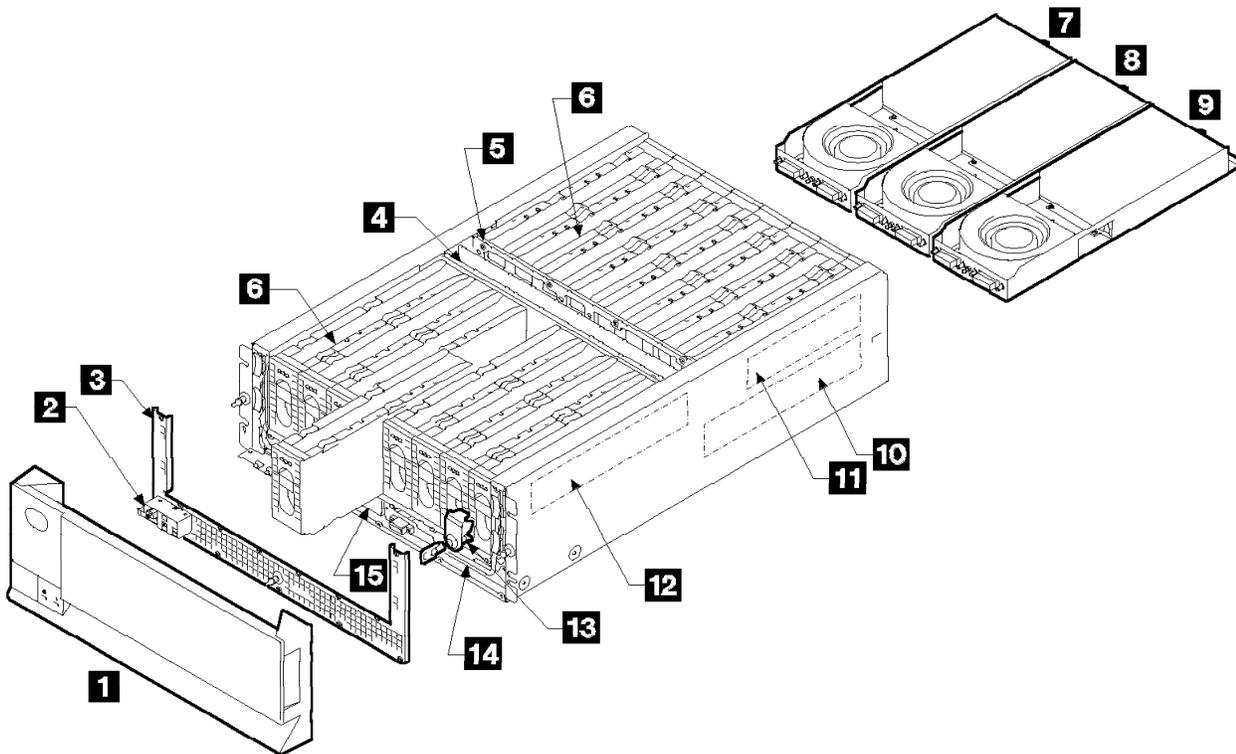


Figure 9. IBM 7133 Component Location. The location of the internal components of the 7133 model 010.

1. Front cover
2. Control panel assembly
3. RFI shield
4. Front backplane assembly

5. Back backplane assembly
6. Disk drive modules
7. Third fan-and-power assembly
8. Second fan-and-power assembly
9. First fan-and-power assembly
10. Front signal card (one on both sides)
11. Back signal card (one on both sides)
12. Back power card (one on both sides)
13. Optional disk drive module lock
14. Right-hand power distribution assembly
15. Left-hand power distribution assembly

Back connectors: The diagram in Figure 10 shows the location and organization of the connectors at the rear of the rack-mounted 7133 model 010.

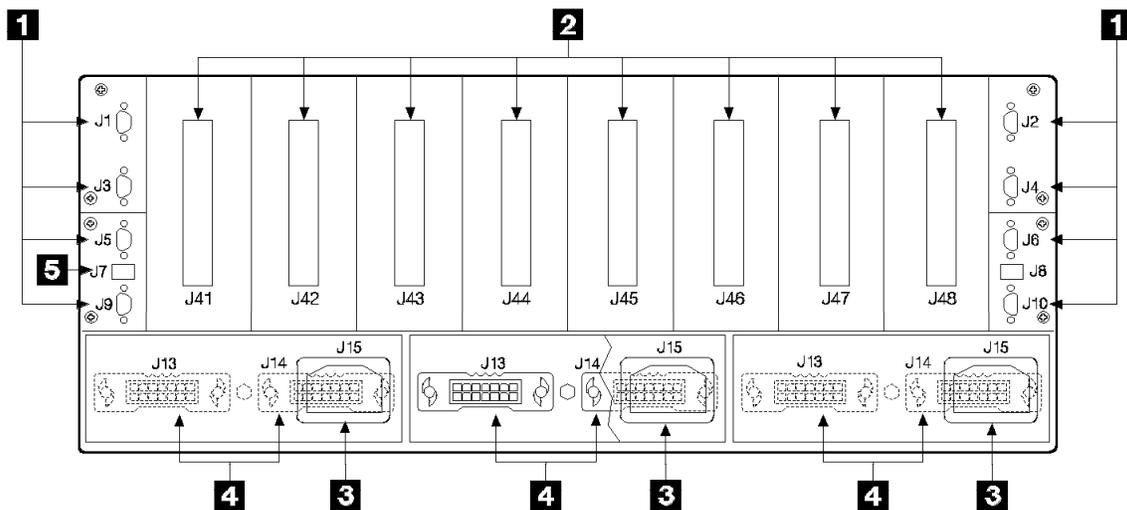


Figure 10. IBM 7133 Back Connectors. The location of the back connectors on the 7133 model 010.

1. SSA connectors
2. Disk drive module connectors
3. Mainline power connectors
4. Fan-and-power-supply assembly connectors
5. Power control connector

Physical Layout of the Model 500: This section will illustrate the location of the various components and indicator lights in the 7133 model 500.

Lights and Power: The diagram in Figure 11 on page 18 shows the location of the indicator lights and power switch in the desktside 7133 model 500.

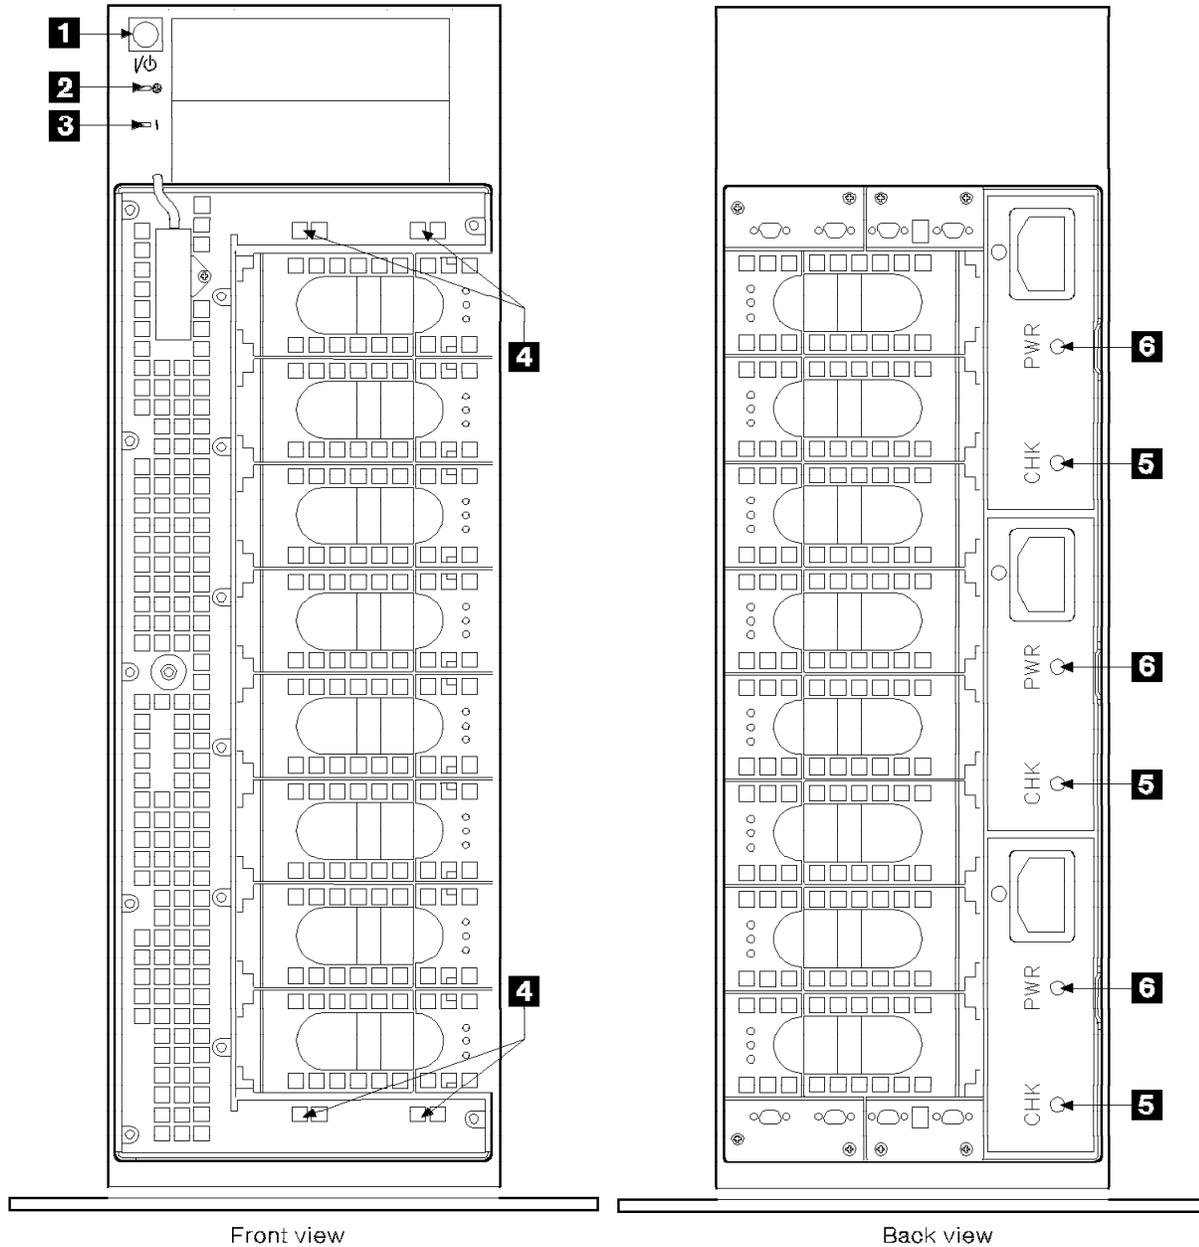


Figure 11. IBM 7133 Lights and Power. The location of the indicator lights and power switch.

1. Power switch

This switch controls the flow of internal DC power from the power and cooling modules to the rest of the unit. Power will always be present in the power and cooling modules if they are connected to an external power source. To activate the 7133, press and hold the switch for two seconds. To deactivate the 7133, again press and hold the switch for two seconds.

2. Subsystem power light

When lit, this green light indicates that the 7133 is being supplied with DC power.

3. Subsystem check light

When lit, this amber light indicates that a failure has been detected within the 7133. The unit may still be able to operate normally, depending upon the nature of the problem.

4. Power card light

When lit, these green lights indicate that DC power is being supplied to the power cards for the disks.

5. Fan-and-power check light

When lit, these amber lights indicate that there is a problem with the DC power supply. If the remote-power-on feature is being used, this light indicates that power has been terminated by the using system or the remote-power-on control unit.

If the light is flashing, it indicates a cooling fan failure.

6. Power light

When lit, these green lights indicate that external AC power is being supplied to the power and cooling modules.

Parts Location: The diagram in Figure 12 shows the location and organization of the various components inside the deskside 7133 model 500.

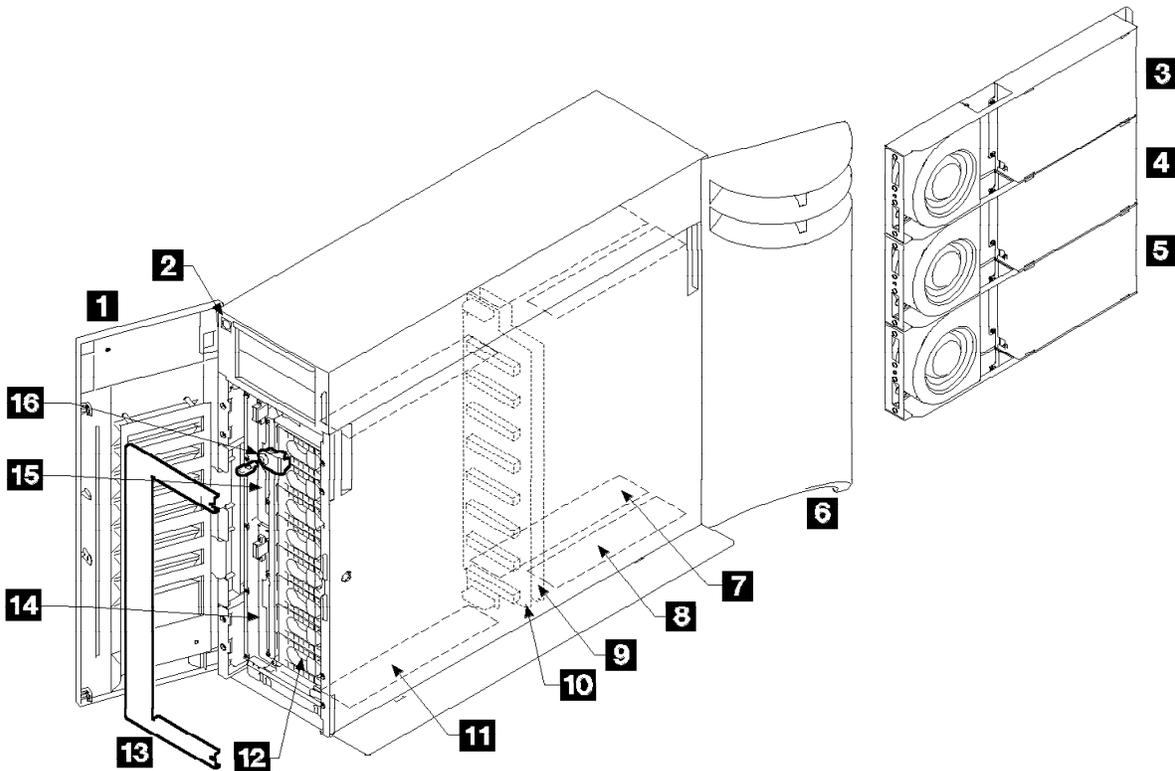


Figure 12. IBM 7133 Component Location. The location of the internal components of the 7133 model 500.

1. Front cover
2. Control panel assembly
3. Third fan-and-power assembly

4. Second fan-and-power assembly
5. First fan-and-power assembly
6. Back cover
7. Front signal card (one at top and bottom)
8. Back signal card (one at top and bottom)
9. Back backplane assembly
10. Front backplane assembly
11. Back power card (one at top and bottom)
12. Disk drive modules
13. RFI shield
14. Lower power distribution assembly
15. Upper power distribution assembly
16. Optional disk drive module lock

Back connectors: The diagram in Figure 13 on page 21 shows the location and organization of the components at the rear of the rack-mounted 7133 model 500.

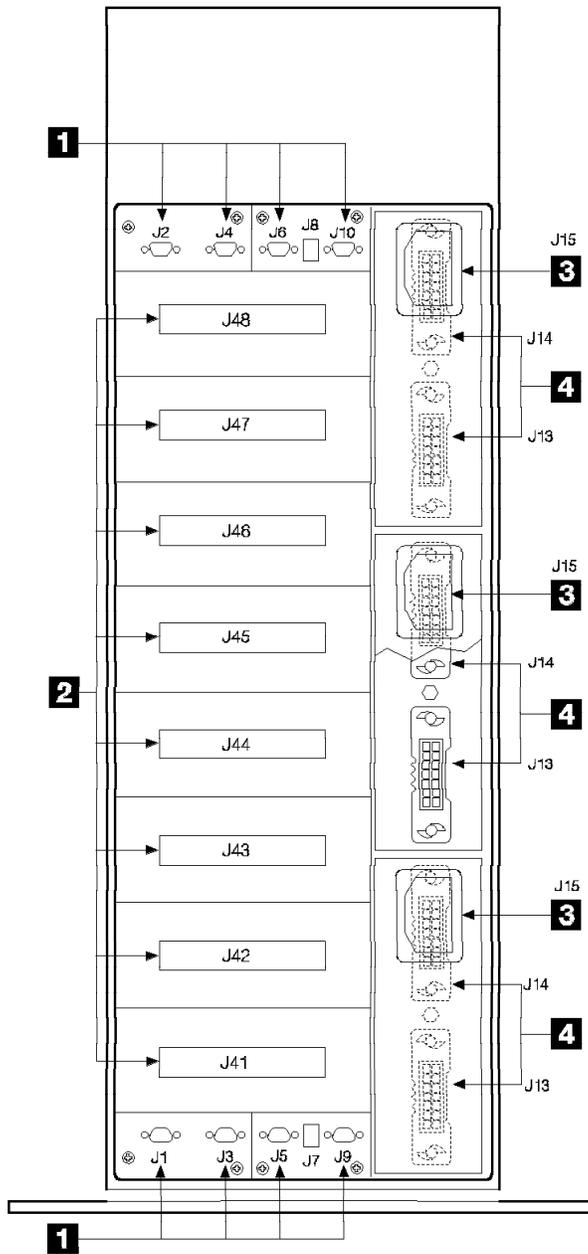


Figure 13. IBM 7133 Back Connectors. The location of the back connectors on the 7133 model 500.

1. SSA connectors
2. Disk drive module connectors
3. Mainline power connectors
4. Fan-and-power-supply assembly connectors

2.1.2.2 Physical Layout of an SSA Disk Drive Module

The 7133 is populated with SSA disk drive modules that have the physical characteristics illustrated in Figure 14.

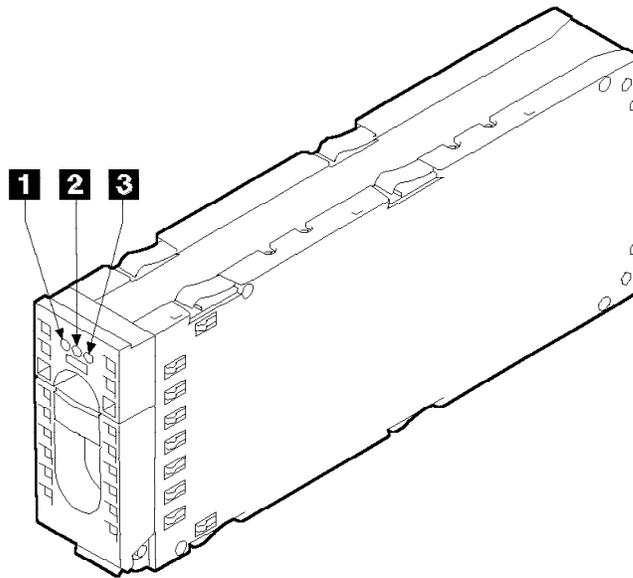


Figure 14. SSA Disk Drive Module. The physical layout of an SSA disk drive module.

1. Power light

When lit, this green light indicates that DC voltage is present and within the specified limits.

2. Ready light

Depending on its state, this green light indicates the following conditions:

<i>Off</i>	Both SSA links are inactive due to one of the following conditions: <ul style="list-style-type: none">• The disk drive modules (dummy or real) that are adjacent to this one are not connected or are missing.• The disk drive modules that are adjacent to this one are inactive.• This disk drive module is running a Power-On Self Test (POST).
------------	--

<i>Permanently on</i>	Both SSA links are active and the drive is ready to accept commands from a using system. This status does not indicate that the disk motor is in operation.
-----------------------	---

<i>Slow flash</i>	The light stays on for two seconds and off for two seconds. Only one SSA link is active.
-------------------	--

<i>Fast flash</i>	The light flashes five times per second. The drive is active with a command in progress.
-------------------	--

3. Check light

Depending on its state, this amber light indicates the following conditions:

<i>Off</i>	Normal operating conditions.
------------	------------------------------

Permanently on One of the following conditions has occurred:

- An unrecoverable hardware error has occurred.
- An unrecoverable error has been detected by the disk drive microcode.
- The Power On Self Tests (POSTs) are running or have failed. The light comes on immediately the drive is powered up and should go out after approximately one minute. If the light is on for longer than one minute, the POST has failed.
- Neither SSA link is active.
- The drive is in service mode and may be removed from the 7133.

Flashing

This state is set by the service aid to allow easy identification of the drive.

2.1.2.3 Physical Layout of the SSA Four-Port Adapter

The 7133 SSA disk subsystem is attached to the SSA four port adapter (feature code 6214, or the newly announced feature code 6216 on the RS/6000). The physical layout of the adapter is illustrated in Figure 15.

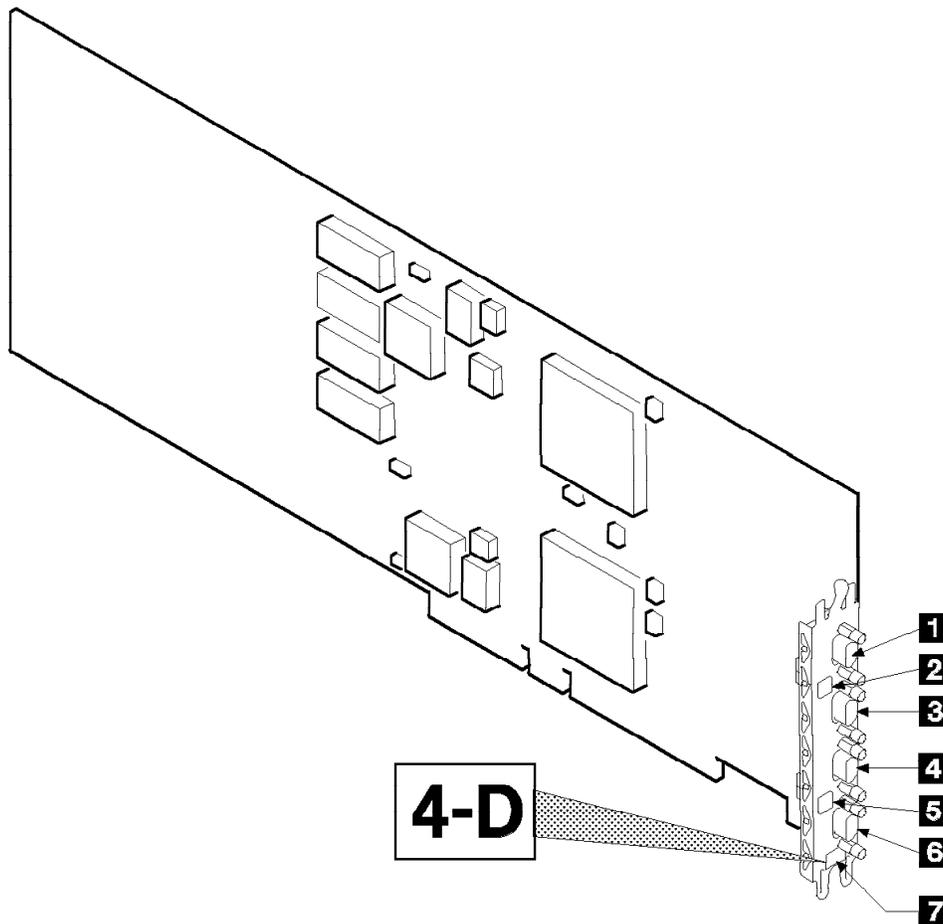


Figure 15. SSA Four-Port Adapter. The physical layout of the SSA four-port adapter.

1. Connector B2
2. Green light for adapter port pair B
3. Connector B1
4. Connector A2
5. Green light for adapter port pair A
6. Connector A1
7. Type-number label

The green lights for each adapter port pair indicate the status of the attached loop as follows:

- Off* Both ports are inactive. If a disk drives are connected to these ports, then either the modules have failed or their SSA links have not been enabled.
- Permanently on* Both ports are active.
- Slow flash* Only one port is active.

2.1.2.4 7133 Hardware Schematics and Connection Rules

A schematic representing the internal layout and connections for the disks in both the model 010 and model 500 is shown in Figure 16. As can be seen, the disks are located in two rows of eight disks within the 7133. Each row is divided into two groups of four disks each, giving four groups in each 7133. Each group is accessible individually via connectors numbered J1 to J10 as shown.

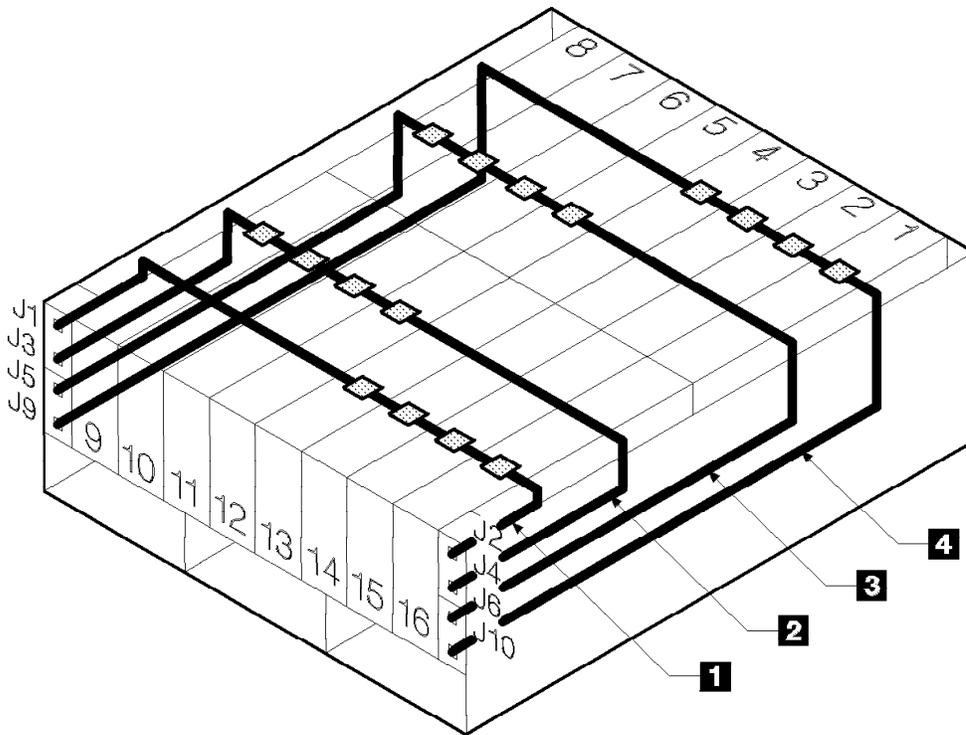


Figure 16. IBM 7133 Hardware Schematic. The layout of the disks and connectors inside the 7133.

In order to simplify the process of identifying disks for operations and problem determination purposes, if it is required to connect strings together, they should

always be connected in a particular way. There are quite a few cables that can be used to connect strings to other strings, subsystems to adapters, subsystems to subsystems, and adapters to adapters though, in actual fact, there is really only one cable; it just varies in length and in part number and if this is borne in mind, connection is in actual fact quite easy. The apparent complexity is only to allow correct tracking and validation of configurations. The following explanation should clarify the rules that must be followed for creating valid SSA configurations.

- Cables

The cables available are listed in Table 2 and in Table 3 on page 26. The first table contains those cables that should be ordered if a 7133-500 or a 7133-010 which is not part of a complete system order (including a host RS/6000), is being configured. Note that identifying cables is easy. All feature codes are four digits long and begin with a 5. The next three digits specify the cable length, 002 for a 0.18 meter cable, 006 for a 0.6 meter cable and 010 for a 1.0 meter cable, for example.

Description	Part	Feature	Length
From 7133 to adapter	07H8985	5010	1.0m
	32H1465	5025	2.5m
	88G6404	5050	5.0m
	32H1466	5100	10m
	88G6406	5250	25m
From group to group	31H7960	5006	0.6m
	07H8985	5010	1.0m
	32H1465	5025	2.5m
	88G6404	5050	5.0m
From group to group (1)	07H9163	5002	0.18m
Note:			
1. This cable is used to connect groups in the same subsystem where the connectors are on the same side of the subsystem.			

The second table contains those cables that should be ordered if a 7133-010 is being configured as part of a complete system order, including a host RS/6000. The cables are the same as in the other table, they just have different part numbers, and the 1.0 meter and 0.6 meter cables are not available.

<i>Table 3. 7133 Model 010 Ordered as Part of a System</i>			
Description	Part	Feature	Length
From 7133 to adapter	32H1465	5904	2.5m
	88G6404	5905	5.0m
	32H1466	5906	10m
	88G6406	5907	25m
From 7133 to 7133	32H1465	5704	2.5m
	88G6404	5705	5.0m
	32H1466	5706	10m
	88G6406	5707	25m
From 7133 to additional adapter	32H1465	5604	2.5m
	88G6404	5605	5.0m
	32H1466	5606	10m
	88G6406	5607	25m
From group to group (1)	07H8985	5803	1.0m
From group to group (2)	07H9163	5801	0.18m
Note:			
<ol style="list-style-type: none"> 1. Do not use a cable length of 0.6 meters between connectors on the same 7133 when the connectors are on opposite sides of the subsystem as the cable interferes with the removal and insertion of disk and power modules. 2. This cable is used to connect groups in the same subsystem where the connectors are on the same side of the subsystem. 			

The purpose and use of each cable is shown in the tables.

- Connecting groups inside a subsystem

For administration purposes, the disks are numbered as shown in Figure 17 on page 27 and must be connected in sequential order. Thus, if group 1 is to be connected to group 2, then disk 4 must be adjacent to disk 5. If group 2 is to be connected to group 3, then disk 8 must be adjacent to disk 9, and so on. This means that when physically connecting the groups inside a subsystem together, connecting group 1 to group 2 would mean using a 0.6 meter cable (feature 5006) inside a 7133-500 or a 1.0 meter cable (feature 5010) inside a separate 7133-010. The 1.0 meter cable is used instead of the 0.6 meter cable to prevent obstruction of the power and cooling modules in the model 010. Therefore a 1.0 meter cable should also be used in a 7133-010 when ordered as part of a system (feature 5803) for connecting group 1 to group 2. This cable is used because in order to connect disk 4 to disk 5, connector J9 must be connected to connector J6 (see the picture in Figure 16 on page 24), and these connectors are on opposite sides of both models of the 7133.

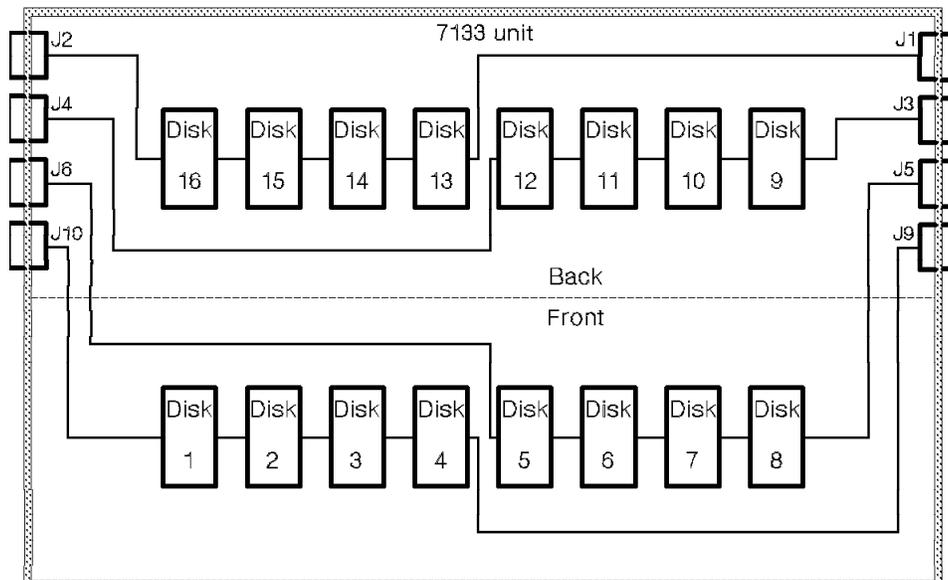


Figure 17. 7133 Disk Numbering. The order in which the disks are numbered internally.

Therefore, if group 2 were to be connected to group 3, disk 8 must be adjacent to disk 9, so connector J5 must be connected to connector J3. This requires the 0.18 meter cable for both models of the 7133. This is feature 5002 for the model 500 and model 010 when ordered separately; it is feature 5801 for the model 010 when ordered as part of a complete system.

It comes as no surprise to hear that connecting group 3 to group 4 therefore requires a 0.6 meter cable for the model 500 and a 1.0 meter cable for the model 010, connecting connectors J4 and J1. For completeness, this would require feature code 5006 for the model 500, feature code 5010 for the model 010 when ordered separately and feature code 5803 for the model 010 when ordered as part of a system.

Please refer to Chapter 4, “Hardware Scenarios” on page 49 for comprehensive examples of valid 7133 configurations.

- Connecting groups between subsystems

If it is required to connect groups of disks in separate subsystems in order to create loops of greater than 16 disks, or just to balance workload, then cables should be selected based upon the distance between the subsystems. When connecting groups together, it does not matter in which order the disks are connected, but it will ease administration greatly and also affect performance if the configuration is planned sensibly. Please refer to Chapter 4, “Hardware Scenarios” on page 49, for information on useful valid configurations and to Chapter 6, “Performance/Availability Investigation” on page 99, for a discussion on how best to plan configurations for performance and availability optimization. The cables defined for subsystem-to-subsystem connections are the 2.5 meter, 5.0 meter, 10.0 meter, and 25.0 meter cables. These particular cables must be used not because they are different in any way but because it will ease administration and problem determination if a consistent set of cables is used.

Referring to Figure 18 on page 28 as an example of how to connect groups in separate subsystems together, it can be seen that group 1 in subsystem 1 has been connected to group 1 in subsystem 2. Disk 4 in subsystem 1 is adjacent to disk 1 in subsystem 2 and a 2.5 meter cable (feature code 5025

for the model 500 and model 010 when separately ordered, feature code 5604 when the model 010 is part of a system) has been used.

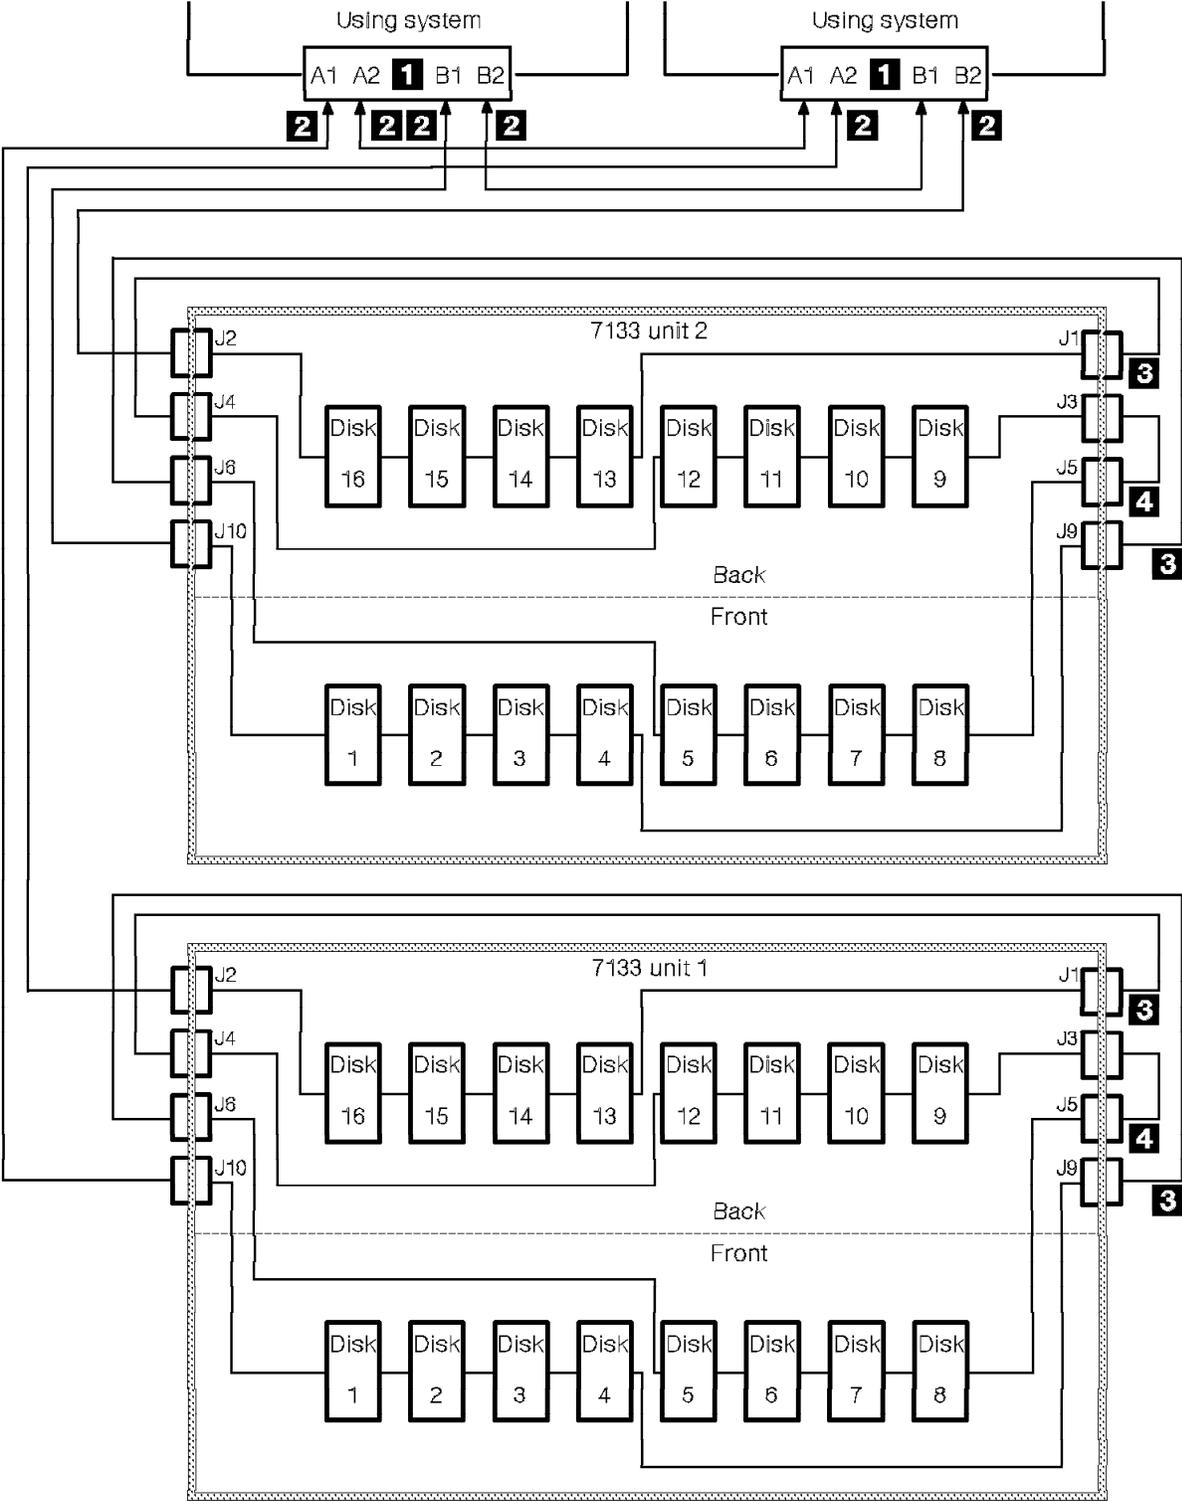


Figure 18. Dual 7133 Configuration. An example of the connections for two 7133 subsystems.

- Connecting groups to adapters

In order to connect a group to an adapter, the cable to use will again depend upon the distance from the group connector to the adapter. Any end of any group can be connected to an adapter though there are limitations on the number of adapters and disks in a loop which are detailed below.

The cables defined for subsystem-to-adapter connection are the 1.0 meter, 2.5 meter, 5.0 meter, 10.0 meter and 25.0 meter cables. A typical example of such a connection is shown in Figure 18 on page 28.

- Creating loops

There are a number of rules which must be followed when creating loops.

1. Maximum adjacent empty slots

In any loop, there can be no more than *three* adjacent empty disk slots, and any empty slot must be filled with a dummy disk drive module. There can be as many empty slots as required (always filled with dummy modules), just never more than three adjacent at any point.

2. Valid adapter port pair

Each loop must be connected to a valid pair of connectors on an adapter card, that is either A1 and A2 or B1 and B2.

3. One pair of ports per adapter

Both pairs of ports on an adapter card can never be in the same loop, each pair must be in a different loop.

4. Two/eight adapter port pairs per loop

A maximum of two adapter port pairs (from different adapters) can be in any one loop when feature code 6214 adapters are used; with the new feature code 6216 SSA adapters, this maximum becomes eight.

5. Only two adapters per host

A maximum of two SSA adapters that are in the same loop can be installed in the same host.

6. Maximum of 48 drives

The maximum number of drives in an SSA loop is 48.

Ensuring that these rules are adhered to, creating a loop is then relatively straightforward. It is vital to re-emphasize that all loops should be well planned and documented. See Chapter 4, "Hardware Scenarios" on page 49, for some valid loop scenarios for many environments and Chapter 6, "Performance/Availability Investigation" on page 99, for information on planning the optimum loop, and Appendix B, "7133 Configuration Planning Charts" on page 161, for some useful planning charts.

Creating a basic loop simply involves connecting one port of an adapter-port pair to a group in a subsystem and then connecting the other end of the group to the other port in the same adapter-port pair. Loops can become as arbitrarily complex as required from this point, within the rules stipulated above. Groups can be connected together and a second adapter-port pair introduced for example. Two examples of relatively simple loops are shown in Figure 18 on page 28.

2.1.2.5 SSA Software Support

The software support for the 7133 SSA Disk Subsystem provides the following components:

- Configuration methods

The configuration methods initiate the configuration of any SSA four-port adapters in the system and of any associated devices connected to the ports of these adapters. Thus, communications with any other adapters and with any disks on connected loops will be established and the ownership of the disks decided so that the necessary configuration can take place.

Device drivers are then created for the adapters and disks on any attached loop and AIX disk device drivers created to interface to these with a one-to-one correspondence. More detail is provided on the device drivers in the next section.

- Device drivers

Four types of device driver are involved:

- Four-port adapter device driver

This device driver supports the four-port adapter providing access to the adapter for communications and for managing the adapter itself. It is normally created at the next system boot after an adapter has been installed.

- SSA router

The SSA router does not actually represent any physical hardware. It exists only to be the parent of SSA logical and physical disk device drivers. The nature of SSA allows disks to be managed by up to two adapters in each host system that the disks are connected to, so that in the event of failure of one of the adapters, the other can take over. For this to occur seamlessly, the disk device drivers are not directly related to a specific adapter, but are able to dynamically switch between adapters.

SSA routers are named `ssar` by default.

- SSA physical disk device driver

This device driver represents the physical properties of SSA disks and is by default named `pdiskn`, where `n` varies for each disk connected (`pdisk0`, `pdisk1`, and so on). One physical disk device driver will be created for each SSA disk connected to the system.

These device drivers are children of the parent SSA router and provide for physical management of their associated disks, including diagnostic functions and the setting of certain physical modes of operation. The physical disk device drivers only provide a character-based interface to the disks, being primarily intended for control functions.

- SSA logical disk device driver

This device driver represents the logical properties of SSA disks and can be included in volume groups and have logical volumes created on it. One logical SSA disk device driver is created for each SSA disk connected and is by default named `hdiskn`, where `n` varies for each disk (`hdisk0`, `hdisk1`, and so on).

These device drivers are also children of the parent SSA router, and provide for logical operations such as disk organization and reading and

writing of blocks of data. The logical device drivers therefore accept character and block reading and writing.

- Diagnostics

The diagnostics software provides the ability to verify the correct operations of the SSA hardware. SSA links and disks can be tested, disks can be certified and formatted and an option to flash a light on individual disks is provided to assist in locating disks for diagnostic and management purposes.

More information on software configuration, management and diagnostics is available in Chapter 3, "Installation and Configuration" on page 33, and in Chapter 5, "Problem Determination and Resolution" on page 77.

2.2 Third-Party Serial Storage Architecture Products

The Serial Storage Architecture standard is not an IBM proprietary standard and is being supported by a number of other manufacturers and vendors. In addition, the standard has been submitted to ANSI which has formed the X3T9.7 committee for the purposes of approving SSA as a standard and defining that standard. As part of this work, a project to map SCSI-3 to SSA has also been accepted by ANSI under the X3T9.2 committee. Currently, the SSA specification is under the control of the Serial Storage Architecture User's Industry Group (SSA-UIG), which was formed in late 1992. This group also has responsibility for the SCSI2-to-SSA mapping.

The following list is not exhaustive, but illustrates the general level of industry commitment to Serial Storage Architecture. The list contains members of the SSA Industry Association and is taken from the Industry Associations World Wide Web page <http://www.ssaia.org/>. Where available, more detail on the members is provided, but only with regard to current SSA support.

- Adaptec
- ADTX
- Alpha Peripherals
- American Kotobuki Electronics
- AMP
- Amphenol
Amphenol provide cables, connectors and terminators.
- C&M Corporation
C&M provide wire and cable products.
- Circuit Assembly Corporation
- Digital Equipment Corporation
- DPT
DPT provides PCI-SSA host adapters and PCI-SSA disk array controllers.
- EBD
- ENDL Associates
- Fujitsu Computer Products
- Future Domain Corporation

- IBM Corporation
 - Microelectronics Division
 - Storage Systems Division
 IBM provides SSA adapters, disks and subsystems.
- Initio
- ITT-Canon
- Madison Cable Corporation
- Micropolis
- Mitsumi Electronics Corporation

Mitsumi provide SSA CD-ROM drives.
- Molex, Inc.

Molex provides SSA cables and connectors.
- Pathlight Technology, Inc.
- Peer Protocols, Inc.
- Samsung Semiconductor
- Siemens Nixdorf
- Silicon Systems, Inc.
- Solution Technology
- Symbios Logic
- Tandberg Data Storage AS
- Tandem Computers
- Vicom
- VLSI Technology
- Western Digital
- Xyratex

Xyratex provides SSA development and production test systems, SSA compliance test software and test services and SSA storage subsystems.
- Zadian Technologies

Zadian provides SSA test systems and SSA development systems.
- Zitel

Zitel designs and develops SSA storage subsystems.

Chapter 3. Installation and Configuration

This chapter will look at installing and configuring the 7133 from both a hardware and software perspective. This will cover both the Model 010 and the Model 500. This chapter should be used in conjunction with:

- 7133 SSA Disk Subsystem: Operator Guide
- 7133 SSA Disk Subsystem: Service Guide
- 7133 Model 010 SSA Disk Subsystem: Installation Guide
- 7133 Model 500 SSA Disk Subsystem: Installation Guide

which should be delivered with the 7133 hardware.

An SSA four-port adaptor (type 4-D, feature number 6214 or 6216) should either be already installed in the using system or available for installation.

3.1 General Procedure

This section looks at the general requirements and procedures for installing and setting up the 7133 disk subsystem.

3.1.1 Software Requirements

The software requirements are:

1. AIX level 3.2.5 Enhancement 5
2. 3251 Preventative Maintenance Package (PMP)
3. AIX level 4.1.4

The following additional Licenced Packaged Products (LPPs) are required for AIX level 3.2.5:

- 5756-03082 ssa
- 5756-03083 ssadisk

For AIX level 4.1.4, the SSA support is included in the following filesets:

- devices.ssa.disk.rte
- devices.mca.8f97.rte
- devices.mca.8f97.com
- devices.mca.8f97.diag

Authorized Program Analysis Report (APAR) IX52018 can be ordered from your local support center, though the hardware should be shipped with all software prerequisites.

To check the level of software use the `lslpp` command as follows:

```

> lspp -h ssa*
Name
-----
  Fix Id  Release          Status   Action   Date       Time       User Name
-----
Path: /usr/lib/objrepos
  ssa.diag.obj
          03.02.0000.0000 COMPLETE  COMMIT   09/08/95   12:47:06 root
  ssa.rte.obj
          03.02.0000.0000 COMPLETE  COMMIT   09/08/95   12:47:07 root
  ssadisk.obj
          03.02.0000.0000 COMPLETE  COMMIT   09/08/95   12:47:08 root
Path: /etc/objrepos
  ssa.diag.obj
          03.02.0000.0000 COMPLETE  COMMIT   09/08/95   12:47:07 root
  ssa.rte.obj
          03.02.0000.0000 COMPLETE  COMMIT   09/08/95   12:47:07 root
  ssadisk.obj
          03.02.0000.0000 COMPLETE  COMMIT   09/08/95   12:47:08 root

```

Figure 19. Listing the SSA Product Software

The next step is to ensure that the subsystem is correctly connected to the host system. Details on the connection can be found later in this chapter.

To obtain information about the ssa adapter card, use the `lscfg` command as follows:

```

> lscfg -vl ssa*
DEVICE          LOCATION          DESCRIPTION
ssa0            00-03            SSA Adapter
  Part Number.....031H7762
  Serial Number.....f5449042
  EC Level.....0000D58332
  Manufacturer.....IBM053
  ROS Level and ID.....1039
  Loadable Microcode Level....00
  Device Driver Level.....00
  Displayable Message.....SSA-ADAPTER
  Device Specific.(Z0).....DRAM=000
  Device Specific.(Z1).....CACHE=0

```

Figure 20. Listing the SSA adapter device information

Where the fields have the following meanings:

- Part Number** Adapter card Field Replaceable Unit (FRU) part number
- Serial Number** Adapter card serial number

EC Level	Adapter card engineering change (EC) level
Manufacturer	Manufacturer and plant code
ROS level and ID	The version of Read-only Storage (ROS) code that is loaded on the adapter.
Loadable Microcode Level	The version of loadable microcode code that is needed for the satisfactory operation of this card.
Device Driver Level	The minimum level of device driver that is needed for this level of card.
Displayable Message	SSA adapter description

3.1.2 Hardware Requirements

The hardware requirements are as follows:

- The SSA 4 port adaptor card.
- For the Model 010 (see Figure 21 on page 36)
 1. Front mounting screws, lock washers, and flat washers
 2. Nut clips
 3. Support rail, left
 4. Rail screws
 5. Back mounting screws and washers
 6. Back clamp plates
 7. Remote-power-on control unit (optional)
 8. Using-system power-control cables (optional)
 9. Support rail, right
 10. Three way power cable
 11. SSA cables
 12. 7133 power-control cable (optional)
 13. Disk drive module locks (optional)

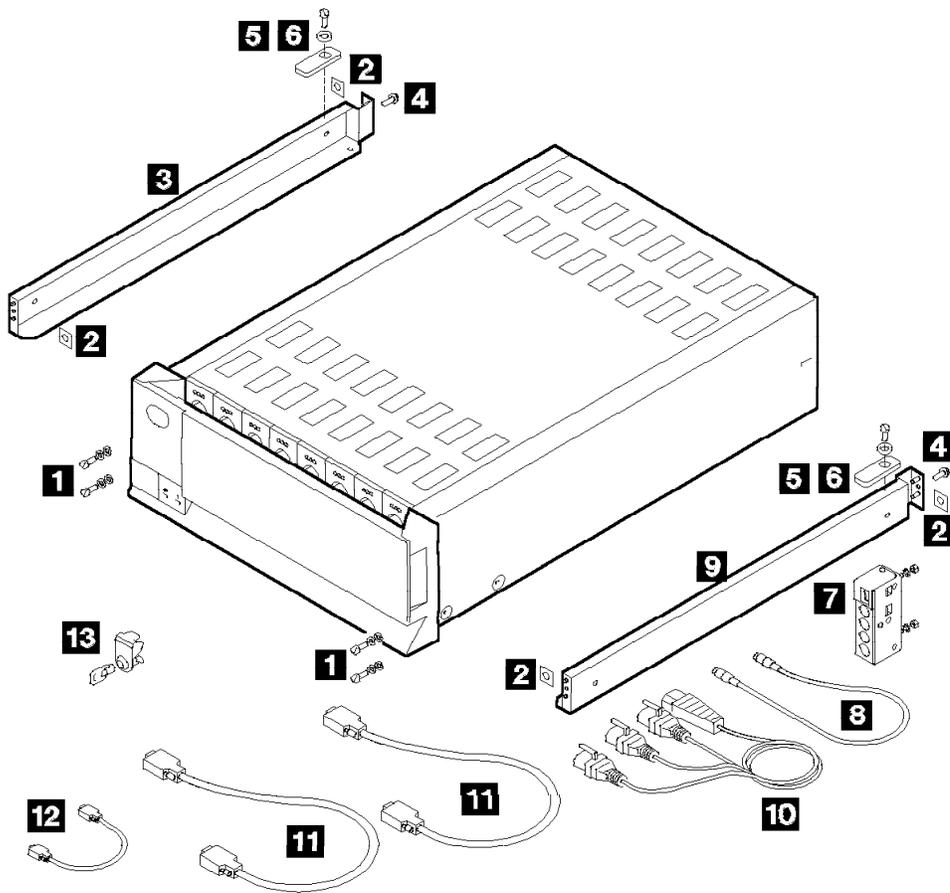


Figure 21. The 7133 Model 010. The parts required to install the 7133 model 010.

- For the model 500 (see Figure 22 on page 37)
 1. 7133 Model 500
 2. Power cable
 3. Three-way converter cable
 4. SSA cables
 5. Disk drive module locks (optional)

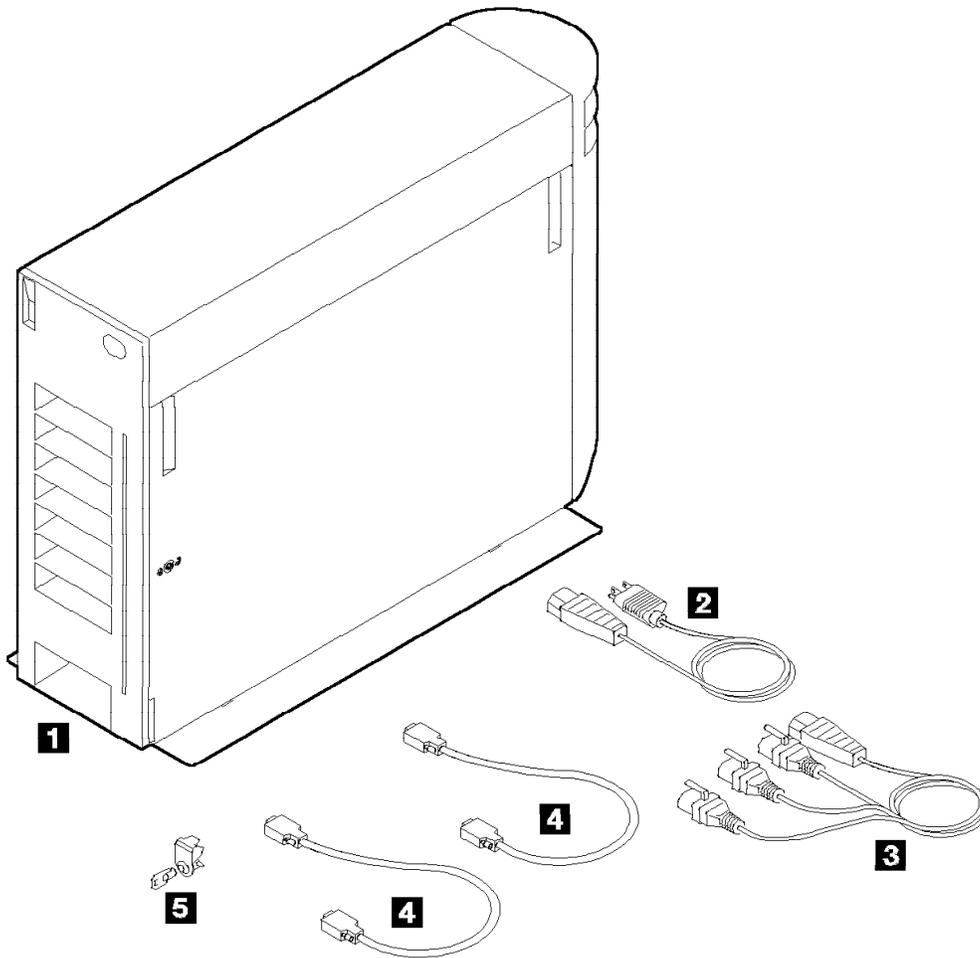


Figure 22. The 7133 Model 500. The parts required to install the 7133 model 500.

Important

The subsystem configuration **must** be known at the time of ordering the 7133 subsystem. The correct length cables and the number of power/cooling modules, disks and dummy modules are dependent on this information. See Chapter 2, “SSA Products Overview” on page 13, for more information on the components to order.

3.2 Installing the SSA Software

The SSA software should be installed on the using host system prior to attaching the 7133 subsystem.

3.2.1 Prerequisites

The following prerequisites must be satisfied prior to installation:

- Back up the system (mksysb).
- Ensure that at least 8 MB of free space are available in the /tmp directory.

3.2.2 Actions

Perform the following actions to install the software:

1. Make sure nobody is currently using the system and that the file system backup is complete.
2. Insert the media containing the ssa and ssadisk products into the relevant device.
3. From SMIT, select the following:
 - **Software Installation & Maintenance**
 - **Install/Update Software**
 - **Install/Update Selectable Software (Custom Install)**
 - **Install software Products at Latest Available Level**

or type `smit install_latest` at the command line.

4. Using the F4 key to list available devices, select the input device.
5. Choose the **SOFTWARE to install** by using F4 to display the list of install images on the media.

Select **3.2.0.0 ssa** and **3.2.0.0 ssadisk**. Toggle the `COMMIT software?` option to no and the `SAVE replaced files?` option to yes. Press **Enter** to begin the installation process. See Figure 23 on page 39 for details.

During the installation process, the screen will display messages while the software is loaded onto the system. If the messages indicate that the installation was not successful, try the process again. Normally, it should work first time. If read or write errors are detected, contact a qualified IBM representative for new install media.

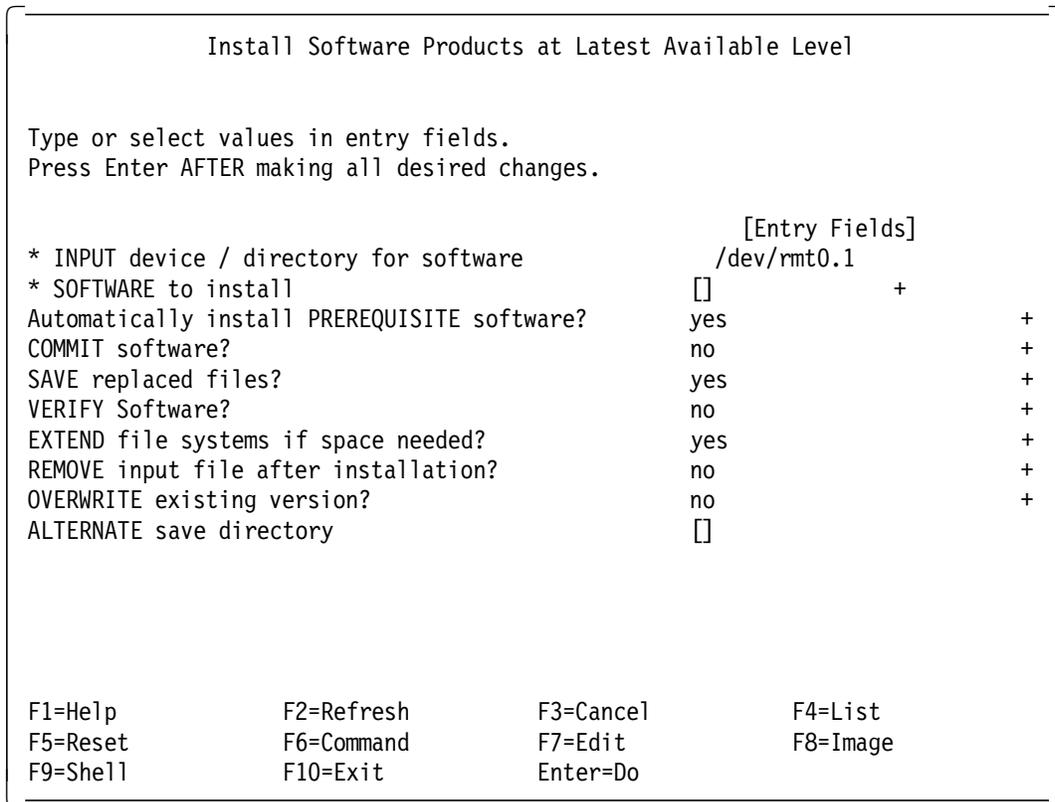


Figure 23. SMIT Software Install Screen

6. Once installation has completed, it will be necessary to shut down the host system. If the four-Port adapter has not yet been installed in the host system, install it while the system is shut down.

3.3 Configuration

At this stage, the subsystem configuration information that was initially used to order the 7133 will be required. Verify that it follows the rules for SSA configurations, as defined in 2.1.2.3, “Physical Layout of the SSA Four-Port Adapter” on page 23.

If the subsystem configuration is not available or valid, it is recommended to use one of the example configurations shown in Chapter 4, “Hardware Scenarios” on page 49.

3.3.1 Basic Configuration

Assuming that the configuration required is a simple one containing only four disk drive modules (see Figure 24 on page 40), these modules *must* be installed in the left-hand four slots at the front of the model 010 and in the top four slots at the front of the model 500. The other 12 slots must contain dummy disk drive modules. The disk drive modules are linked in groups of four; each group of four modules should be linked to two external SSA connectors for a balanced system, but this is not compulsory, as can be seen in Chapter 4, “Hardware Scenarios” on page 49.

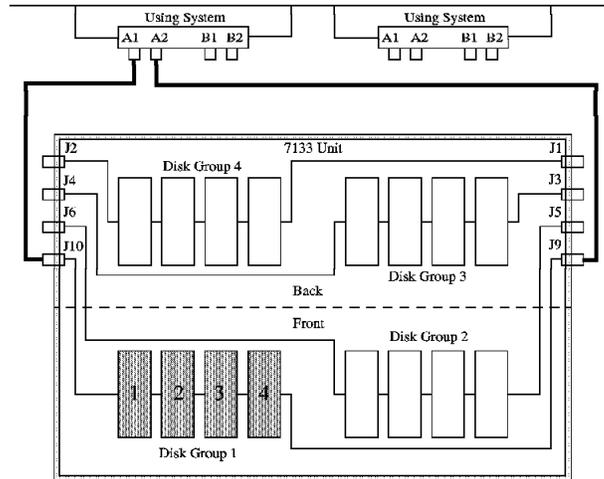


Figure 24. A Basic 7133 Configuration. An example of a simple four disk 7133 configuration

3.3.1.1 Basic Configuration of the 7133

This section deals with setting up the 7133 with a basic four disk configuration. This covers both the Model 010 and Model 500 and assumes the following:

- The adaptor has been installed in the using system.
- The software has been installed on the using system.
- The using system is up and running.
- For the Model 010:
 - The rack has already been set up.
 - The 7133 frame has been installed into the rack.
 - The disks are located in slots 1 to 4 (front left-hand side) of the 7133 frame.
- For the Model 500:
 - Preinstallation planning has been carried out and the tower has been installed correctly with the disks located in slots 1 to 4 (front and top).
- A subsystem-to-adaptor cable from port A1 on the adapter is attached to connector J10 on the 7133.
- A subsystem-to-adaptor cable from port A2 on the adapter is attached to connector J9 on the 7133.
- The SSA adaptor has automatically configured the SSA link and all power and ready lights are on, with all check lights off.

For further information, see the *7133 SSA Disk Subsystem Installation Guide*.

1. Type `cfgmgr` at the command line of the host system to configure the disks into AIX. This should define *hdisks* and *pdisks* to the system. By default, for every physical disk in the subsystem, one *hdisk* and one *pdisk* will be created. Use `lsdev -Cc disk` and `lsdev -Cc pdisk` to list the devices created.

Note.

The relationship between pdisks and hdisks cannot be determined using this command. Use SSA Service Aid - Configuration Verification to establish this relationship.

2. Once this has completed, the configuration should be verified by using the Link Verification service aid as follows:

Type diag on the host system. When prompted, press **Enter** to continue, move the cursor to Service Aids and then to SSA Service Aids. This should produce the following screen:

```
SSA SERVICE AIDS                                     80238

Move cursor onto selection, then press <Enter>.

Set Service Mode
Link Verification
Configuration Verification
Format Disk
Certify Disk

F3=Cancel          F10=Exit
```

Figure 25. SSA Service Aids Menu

- a. Select **Link Verification** and then select the relevant adapter (if there is more than one installed). A list showing the status of all the disk drive modules that are attached to the adapter will be displayed:

```

LINK VERIFICATION 802386

SSA Link Verification for:
  ssa0      00-03   SSA Adapter

To Set or Reset Identify, move cursor onto selection, then press <Enter>

Physical      Serial#      Adapter Port      Status
              A1  A2  B1  B2
-----
pdisk0        11111111    0   3                Good
pdisk1        22222222    1   2                Good
pdisk2        33333333    2   1                Good
pdisk3        44444444    3   0                Good

F3=Cancel      F10=Exit

```

Figure 26. SSA Link Verification Screen

- b. If the output is similar to that shown in Figure 26, then the disks are correctly configured within AIX; otherwise, consult the *7133 Service Guide* and Chapter 5, “Problem Determination and Resolution” on page 77.
3. Go back to the Diagnostic Routines menu within diagnostics and select the **System Verification** menu. Select each of the pdisks listed in turn to certify the disks.
4. If this all completes successfully, then the SSA disks are ready for use; otherwise, use the *7133 Service Guide* to determine the cause of the problem.

3.4 Additional Configurations

Once the basic configuration has been set up, changes can easily be made to it, such as adding or removing disks or moving disks around within a loop.

3.4.1 Adding Disks

As mentioned in Chapter 1, “SSA Technology Overview” on page 1, SSA disks are hot pluggable, which means that neither the host system nor the 7133 subsystem needs to be powered down when adding or removing disks. The only thing to remember is that any new configuration *must* follow the rules laid out in Table 1 on page 14 and 2.1.2.4, “7133 Hardware Schematics and Connection Rules” on page 24.

So, for example, if it was required to add one to four more disks to the basic configuration, here are some of the things that need to be considered:

- The additional disks should go in slots 5 to 8 sequentially, so for adding one disk, it goes in slot 5. For two disks, they go in slots 5 and 6 and so on. All the unused slots within a bay should contain dummy modules.
- Either the additional disk(s) can be added to the current loop, or they can be set up as a new loop. For performance reasons, creating an additional loop may be recommended.

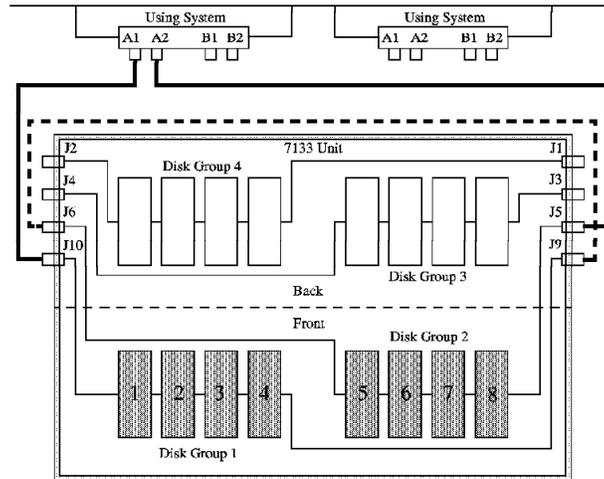


Figure 27. 8 Disk Configuration. All eight disks in one loop.

The configuration of the SSA cables if adding the additional disks to the current loop is as follows:

- A1 to J10
- J9 to J6 (round the back of the 7133)
- J5 to A2

See Figure 27.

For a balanced load across two loops:

- A1 to J10
- J9 to A2
- B1 to J6
- J5 to B2

See Figure 28 on page 44.

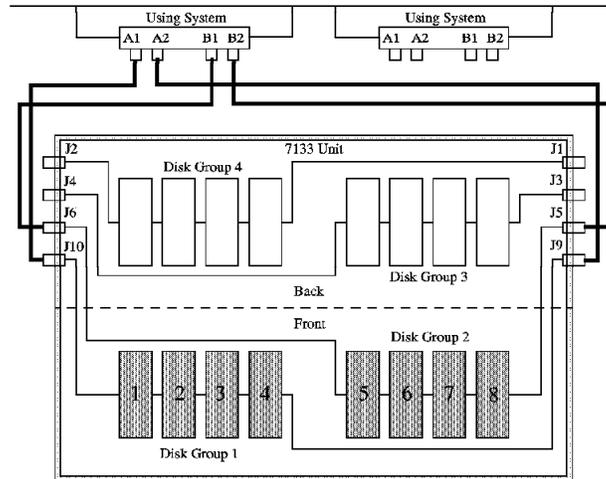


Figure 28. Eight Disk Configuration. Two loops, balanced.

- With the basic configuration of four disks, two power/cooling units provide redundancy to the subsystem; with five to eight disk drives, a third power/cooling unit is required for redundancy.

Once the disk(s) have been added to the system, running the `cfgmgr` command should configure the additional disks into the system. Use the `lsdev` command and the SSA Service Aids to verify this (see 3.3.1.1, “Basic Configuration of the 7133” on page 40).

3.4.2 Removing Disks

Again, the process is as simple as for adding disks; remember that if a disk is being *physically* removed from the subsystem, then it should be replaced by either another disk or a dummy module.

The process for removing a disk is as follows:

1. Remove all data from the disk, (`smit migratepv`, for example)
2. Remove the disk from the volume group (`smit reducevg`)
3. Remove the hdisk reference from the system configuration (`rmdev -l <hdisk> -d`)
4. Execute the `diag` command, and then select **Service Aids** and then **SSA Service Aids**. Select **Set Service Mode** and then move the cursor to the appropriate `pdisk`. Press **Enter**, and then select the **Set or Reset Service Mode**:

```
SET SERVICE MODE802382

pdisk1    5AEA10C8    00-03-P    2GB SSA C Physical Disk Drive

Move cursor onto selection, then press <Enter>.

+ Set or Reset Identify.
  Select this option to set or reset the Identify indicator
  on the disk drive.
> Set or Reset Service Mode.
  Select this option to set or reset Service Mode on the
  disk drive.
  ENSURE THAT NO OTHER HOST SYSTEM IS USING THIS DISK DRIVE
  BEFORE SELECTING THIS OPTION.

F3=Cancel      F10=Exit
```

Figure 29. Setting Service Mode

Having selected the **Set or Reset Service Mode**, the original SET SERVICE MODE screen will be displayed. The > against pdisk1 shows that this disk is now in service mode; it's amber check LED should be lit. The disk module can now be removed from the system. Pressing **F3** now will produce the following screen:

```

SET SERVICE MODE 802381

Move cursor onto selection, then press <Enter>

pdisk0    5AEA1061  00-03-P  4GB SSA C Physical Disk Drive
> pdisk1  5AEA10C8  00-03-P  2GB SSA C Physical Disk Drive

-----
If you leave this service aid, Service Mode and
Identify will be reset.

Press F3 to remain in the Service Mode menu.

Press F10 to exit the SSA Service Aid and return
to the Service Aids Selection menu

Press <Enter> to leave this Service Aid and reset
Service Mode and Identify.

F3=Cancel      F10=Exit      Enter
-----
F3=Cancel

```

Figure 30. Leaving Set Service Mode

Press **Enter**, and this should redisplay the SSA Service Aids main menu.

- Exit diagnostics, and remove the pdisk reference from the Object Data Manager (ODM) by typing the following command at the command line: `rmdev -l <pdisk_number> -d.`

3.4.3 Moving Disks

Moving disks is also very easy to accomplish with SSA. For safety, the procedure for removing a disk should be followed as described in 3.4.2, "Removing Disks" on page 44.

The disk can then be re-inserted somewhere else on the loop, or even in another loop entirely, but it will retain its SSA serial number, and if it is still accessible by the same host system, its pdisk-to-hdisk mapping. Once allocated, the pdisk information for a disk remains constant regardless of where it is moved to within the system.

Again, for safety, it is best to use the procedure for adding the disk described in 3.4.1, "Adding Disks" on page 42, and then restore any information that was backed up.

3.4.4 Changing the Path to a Disk

It is possible to have up to two adapters on a loop when those adapters are in the same host system; up to eight adapters in total are supported on a loop with the new SSA 4 port adapter (feature code 6216). In the event that a second adapter is present on the loop, each disk will be assigned a primary adapter and a secondary adapter. This static configuration is performed once, and the adapter closest to each disk is assigned as the primary.

The primary adapter can be changed using the `chdev` command as follows:

```
> chdev -l pdiskn -a "primary_adapter=<name>"  
>
```

Where `n` is the number of the `pdisk` to be changed and `<name>` can be:

- adapter_a* Sets the primary adapter to be the value of the `adapter_a` attribute.
- adapter_b* Sets the primary adapter to be the value of the `adapter_b` attribute.
- assign* When the device is made available, the system will calculate the shortest path to this device and assign that adapter to be the primary adapter.

By default, the primary adapter is the shortest path to a disk, and the secondary will be the other route available. It might be advantageous to change this to improve performance in a two adapter scenario.

Note

The primary adapter attribute is only meaningful when both adapters involved are in the same system. The primary adapter attribute has no effect if the adapters are in different systems. Changing the path to a disk may be of benefit in availability scenarios; please see 4.2, "Availability Related Scenarios" on page 60 for more information.

3.5 Administration of the subsystem

Further administration of the subsystem involves the following operations:

- Identifying physical disks in a subsystem
- Link verification
- Configuration verification
- Disk certification
- Disk formatting
- Locating disks without the service aids

and these are described in more detail in Chapter 5, "Problem Determination and Resolution" on page 77.

Chapter 4. Hardware Scenarios

This chapter will investigate some of the possible SSA hardware configurations. The focus will be on providing information on how best to lay out the required storage capacity from the perspective of performance and availability. Therefore, performance and availability considerations will be discussed, and then several example implementations will be given for specific environments.

The final section will look at the situation where both performance and availability are important and the trade-offs that need to be made.

This chapter will focus on the hardware implementation; more information on tuning the software for performance and availability can be found in Chapter 6, "Performance/Availability Investigation" on page 99.

4.1 Performance-Related Scenarios

This section will look first of all at the criteria that must be considered in order to design an SSA solution. Once the requirements have been evaluated against these criteria, a solution can be arrived at, and the second part of this section presents several scenarios to illustrate the process.

4.1.1 Considerations

There are many elements that must be considered in designing an SSA solution for performance, and the purpose of this section is to examine the most important and illustrate their impact on the final performance of an SSA subsystem.

4.1.1.1 Bandwidth

The bandwidth of a communications link governs the maximum speed at which data can be transmitted over the link. The figure that is given is usually the maximum peak or burst speed and not that which can be sustained for any length of time. Furthermore, it should be noted that each component of a system may have different bandwidths and there will be bottlenecks. The figure used should be the one defining the performance of the slowest component in the system and should be the sustained data transmission bandwidth.

For example, SCSI-2 fast and wide defines a bandwidth of 20 MB per second for the bus, but in reality an application will never see this kind of bandwidth.

Looking at the components of the bus:

- Disks

The latest technology SCSI disks currently available have peak transfer rates of around 7 MB per second though the actual sustained rate will be lower due to various disk latencies that must be taken into account. As with SSA disks, throughput is maximized with large block size sequential I/O, and applications with this kind of profile will attain throughput close to the maximum bandwidth.

- Bus

The bus itself has a peak rate of 20 MB per second, but there will usually be more than one device on the bus competing for use. As only one device can

be active at a time, the bandwidth must be shared between devices, and the effective bandwidth will be less.

- Adapter

The adapter is also a bottleneck. Actual transmission speeds through the adapter may be a lot less than the bus speed.

It can be seen, therefore, that the effective bandwidth depends on a number of factors, some controllable and some not. For SSA, the situation is slightly different:

- Disks

SSA disks currently have a peak transfer rate of about 7 MB per second. Though this is a peak figure, the actual rate will depend upon the nature of the application reading and writing data. Sustained sequential I/O will approach this figure.

- Bus

There is no bus in SSA; each disk is connected to the next in a serial chain, where each connection is full duplex operating at 20 MB per second. This gives an effective bandwidth of 40 MB per second for each link and as each disk has two links, an overall potential bandwidth of 80 MB per second. However, since communications must be passed disk to disk, and due to various queuing latencies, the effective sustained bandwidth may be slightly lower.

Also different to SCSI is the fact that there is less competition for the communications link. Each device receives a fair and equal opportunity to communicate. Features like spacial reuse and wormhole routing which are described in Chapter 1, "SSA Technology Overview" on page 1, also help to allow higher bandwidths to be sustained.

- Adapter

The adapter will still be a bottleneck since all conversations must currently originate from an adapter. In addition, the performance of the MicroChannel limits the throughput to the adapter giving a realistic maximum of about 35 MB per second. If this does become an issue, it would be possible to alleviate this problem somewhat by using multiple adapters, and this is described in 4.1.1.3, "Multiple Adapter Access" on page 52.

The actual volume of data that reaches the bus from applications does depend on many factors though and to maximize the use of the available bandwidth, the total volume of data on a loop should not exceed the maximum bandwidth of the loop. However, the maximum volume of data that can be presented will depend on many other factors including CPU speed, execution path length, queuing latencies and workload composition, for example. It is therefore necessary to fix some of these variables and observe the maximum bandwidth achievable within the limitations of the others. The main element that can be fixed is the workload composition. Workload compositions vary enormously from all sequential reads through to all random writes, with every combination in between. High block size sequential writes will generate the highest volume of data on the loop and small block size random writes will generate the highest number of transactions. Throughput in each of these boundary cases should be considered as all other workload compositions will fall between these two extremes.

In order to determine the optimum design for a given workload, it is really necessary to run a benchmark; furthermore, it may be that the design criteria include other considerations, such as cost. In any case, from actual performance measurements and from SSA design theory, it can be seen that in the absence of other factors, to utilize the full bandwidth of a link the I/O profile would have to be 50 percent reads and 50 percent writes. Therefore, for optimum performance, a maximum of eight disks should be on any one loop, given that I/O to each disk is 50 percent reads and 50 percent writes, and that each disk is being driven at its maximum throughput. This maximum number will change with the I/O profile. For example, if an application is doing 100 percent writes then the knee of the performance curve will be at four disks. It can also be seen that limiting factors such as CPU speed and MicroChannel usage will prevent the attainment of much greater throughput than 35 MB per second on a single system.

It is also important to stress, that this boundary case, where an application is doing nothing more than generating 50 percent reads and 50 percent writes of large block size sequential I/O, is unlikely to reflect the actual performance characteristics of any real application. The design should therefore ensure that the total throughput from the applications using a loop should not exceed 35 MB per second for that loop. This could be 35 disks all being used simultaneously for sequential I/O at 1 MB per second to each disk, which is perfectly acceptable if 1 MB per second per disk is the maximum data rate from the applications.

4.1.1.2 Data Path

The data path essentially describes the route from the adapter to the disk that data is destined for. The logic is fairly simple in that, as with most things, the quickest route should be the shortest. Thus, the effective bandwidth for data being transferred to the first disk in a loop should be slightly higher than that for data transferred to the second, which should be slightly higher than that for data transferred to the third, and so on.

However, in practice, the effect of this is minimal and is overshadowed by other factors:

- Wormhole routing

This mechanism optimizes the transfer of data through a device that the data is not addressed to.

- Satisfaction

The satisfaction token guarantees that each disk has a fair chance to send its own data, thereby efficiently allowing the bandwidth to be shared.

- Routing

The route taken to get to a disk can be controlled. For example, in a two adapter scenario, such as that shown in Figure 31 on page 52, the first adapter will own disks 1, 2, 5, and 6, while the second adapter will own disks 3, 4, 7, and 8. This is the default, but can be altered. Generally speaking, the default case will be the optimum in terms of routing. As can be seen, this ownership ensures a maximum of two hops to any disk.

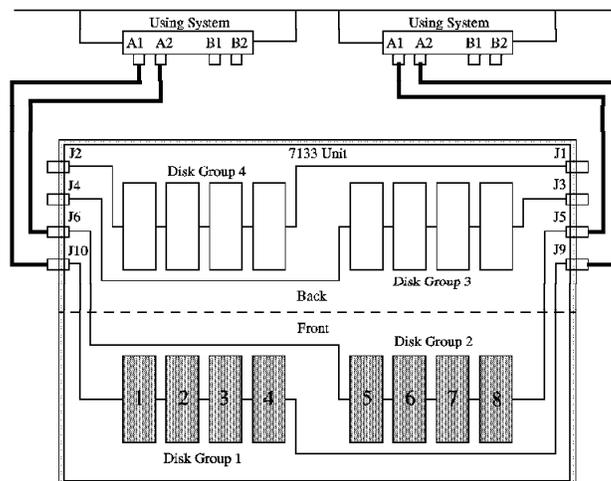


Figure 31. Disk Ownership Example. This illustrates the division of disk ownership.

The routing can be controlled and there may be circumstances in which a more optimum routing can be manually effected. Generally speaking, it is sensible to try and arrange the physical organization of data so that access is uniform across the disks in a loop. This is mainly common sense since if one disk is operating at 100 percent utilization while other disks in a loop are under-utilized, then in all probability, throughput will not be optimum. Rebalancing the load so that all disks are equally utilized will optimize throughput. Furthermore, the way in which the fairness algorithm is implemented means that at maximum utilizations only, the disks furthest from the adapter will receive slightly better service.

4.1.1.3 Multiple Adapter Access

SSA allows disks to be shared between up to eight adapters (when using feature code 6216 SSA adapters; up to two adapters when using the feature code 6214 adapters) so that in the event of failure of an adapter, an alternate path to the disk exists. Up to two adapters can in any one system, so for the maximum of eight adapters, four host systems would be required. This type of scenario presents several possibilities in terms of organization.

Generally speaking, the adapter will be the bottleneck, being limited by factors such as MicroChannel performance to 35 MB per second. At maximum throughputs to disks, and with an I/O profile of 50 percent reads and 50 percent writes, then the maximum number of disks that can be on a loop is eight. Again this is wholly dependent on the applications driving the workload, if the maximum throughput to each disk is 2 MB per second then around 16 disks can be attached to the loop before saturation of the bandwidth. However if the I/O profile is small block size random I/O, with a transaction rate of 2000 I/Os per second then at least 26 disks will be required on the loop to handle the transaction load assuming that data can be evenly spread across them and a maximum transaction rate of 70 I/Os per second per disk; this load would be well within the transaction performance limit of the adapter.

In terms of whereabouts in the loop the adapters should be, optimum performance with increased availability is possible with a configuration such as the one in Figure 32 on page 53. This illustrates a setup with two loops of eight

disks split between two adapters. Again, assuming the optimum I/O profile, an increased bandwidth of around 45 MB per second can be achieved at the extra cost of the second adapter. This scenario is also described from an availability standpoint in 4.2, "Availability Related Scenarios" on page 60.

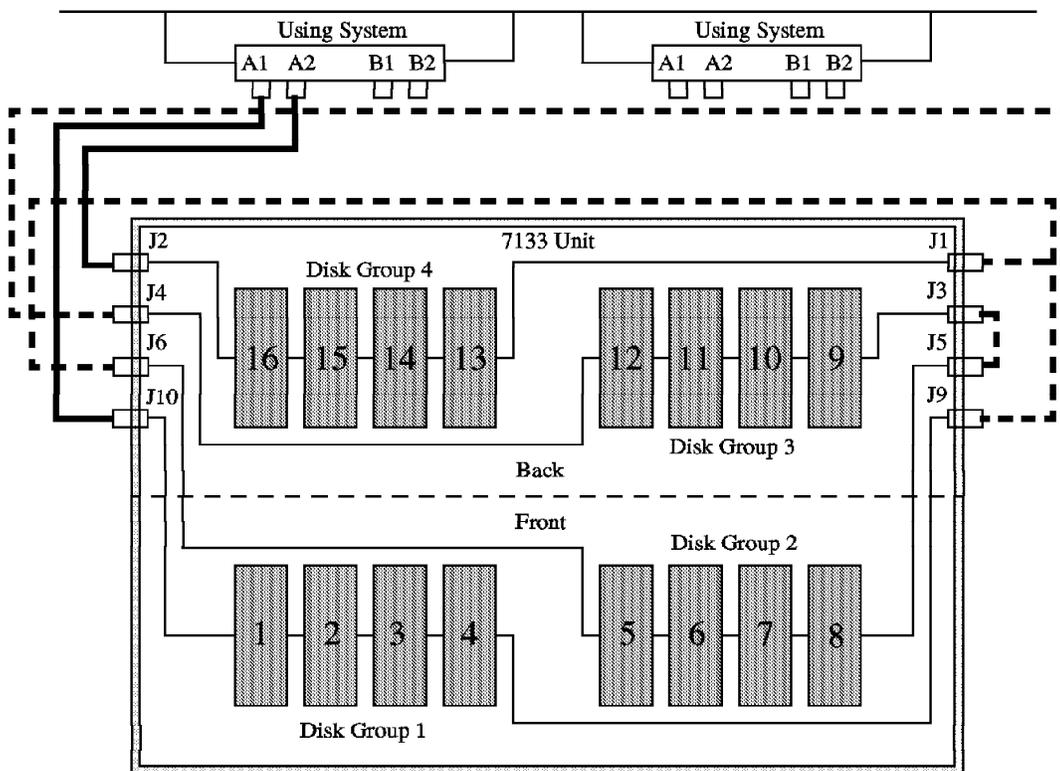


Figure 32. Multiple Adapters Example. An example configuration involving multiple adapters.

4.1.1.4 Data Location

The physical location of data on the disk being accessed will also affect the performance. Traditionally, data located in the center of a disk is statistically more likely to be accessed faster. This is because, on average, the disk head is over the center more often than anywhere else. Control over the location of data on an individual disk can be effected from the Logical Volume Manager, and given that there is an even pattern of access to the disk, statistics show that data in the center of the disk will exhibit better performance.

Thus, performance can be improved by placing the most frequently accessed information at the center of disks and by ensuring that data is uniformly spread across individual disks and between disks in a loop.

It should also be noted that optimum throughput is attained for data located at the outside part of the disk. This is due to the way in which the data is banded on the disk and means that for instantaneous performance, optimum location is attained for data at the outer edges of the disks.

4.1.1.5 Multiple Disk Types

SSA allows loops to be created with mixes of any of the currently available SSA disks. A loop could thus be created with a 4.5 GB, a 2.2 GB and a 1.1 GB disk in it. The mix of disks attached in a loop will not directly affect the overall performance of the loop though it will affect the actual design of the solution. If the overall capacity requirement was 20 GB for example, then several different design possibilities exist. Two examples would be choosing 10 x 2.2 GB disks or choosing 4 x 4.5 GB disks.

From a cost perspective, the 4.5 GB solution may be cheaper since only a single loop and fewer disks are required. If it is possible to spread the data evenly across the 10 x 2.2 GB disks in two loops, performance may be better if the application is throughput limited. It is also now more likely that data can be read concurrently from more disks, thereby increasing effective bandwidth.

4.1.1.6 Higher-Level Requirements

The other factors which must of course be taken into consideration are the higher-level data requirements. This means looking at how many logical volumes and volume groups, for example, are necessary. If many volume groups are required, then at least one disk per volume group will be required which will of course affect the design decisions.

4.1.2 Scenarios

This section will illustrate the design process using several specific scenarios and showing how the above points can be considered to develop an optimum performance solution. In each scenario, the storage requirements will be defined with regard to performance objectives, and then solutions will be defined for various storage capacities.

Note that in the diagrams, the following points relate to the disks:

- No number
There is nothing in this disk slot.
- A number but no shading
There is a dummy disk unit in this disk slot.
- A number and shading
There is a disk in this slot.

4.1.2.1 Scenario 1

In this scenario, analysis of existing I/O information shows a large block size sequential access to large data files. I/O is 50 percent reads and 50 percent writes.

Storage Capacity Requirements - 20 GB: This can be configured as shown in Figure 33 on page 55.

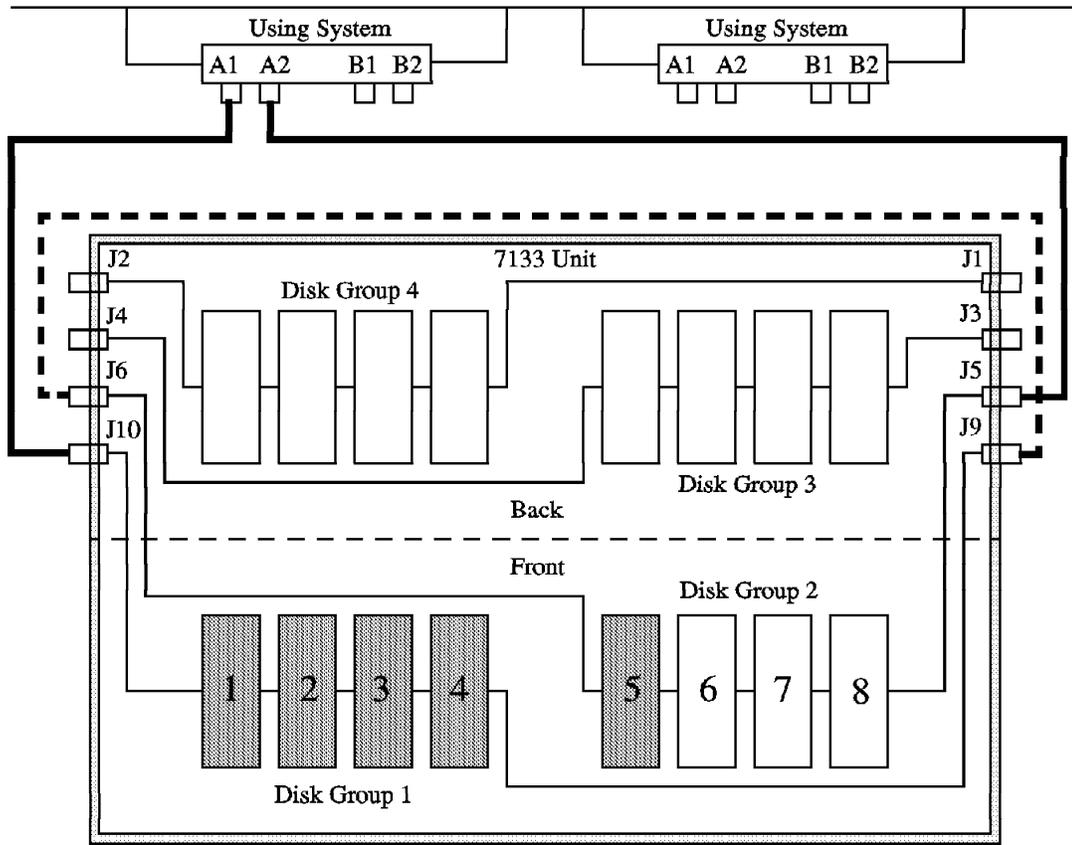


Figure 33. Sequential Access, Capacity 20 GB. This scenario illustrates a design for access to 20GB.

As will be apparent from the discussions in this redbook, there are many configuration possibilities for every requirement. In this particular case the applications are able to drive each disk at 5 MB per second meaning that with 5 disks, the bandwidth of the loop is not likely to be exceeded. The transaction rate is low since the I/Os involve large blocks, so the optimum design would be to use the 4.5 GB disks as shown.

Storage Capacity Requirements - 60 GB: This can be configured as shown in Figure 34 on page 56.

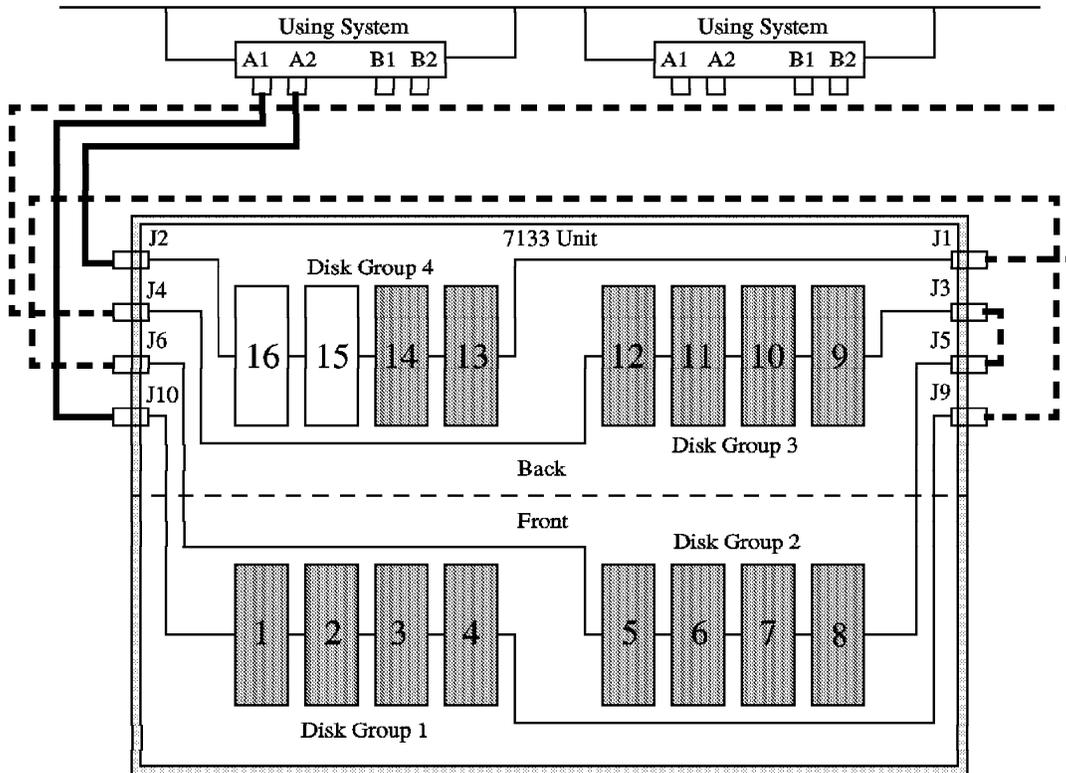


Figure 34. Sequential Access, Capacity 60 GB. This scenario illustrates a design for access to 60 GB.

Once again, the data access will be uniform and will consist of I/O to large evenly distributed files. The applications generating the I/O to the files are capable of sustaining 2 MB per second throughput at peak times. This can be configured in a number of ways, but since cost is an issue here, the cheapest solution will involve a loop of 14 4.5 GB disks.

Once again, the transaction limits of the adapter and disks are not in danger of being exceeded due to the nature of the I/O load.

Storage Capacity Requirements - 120 GB: This can be configured as shown in Figure 35 on page 57

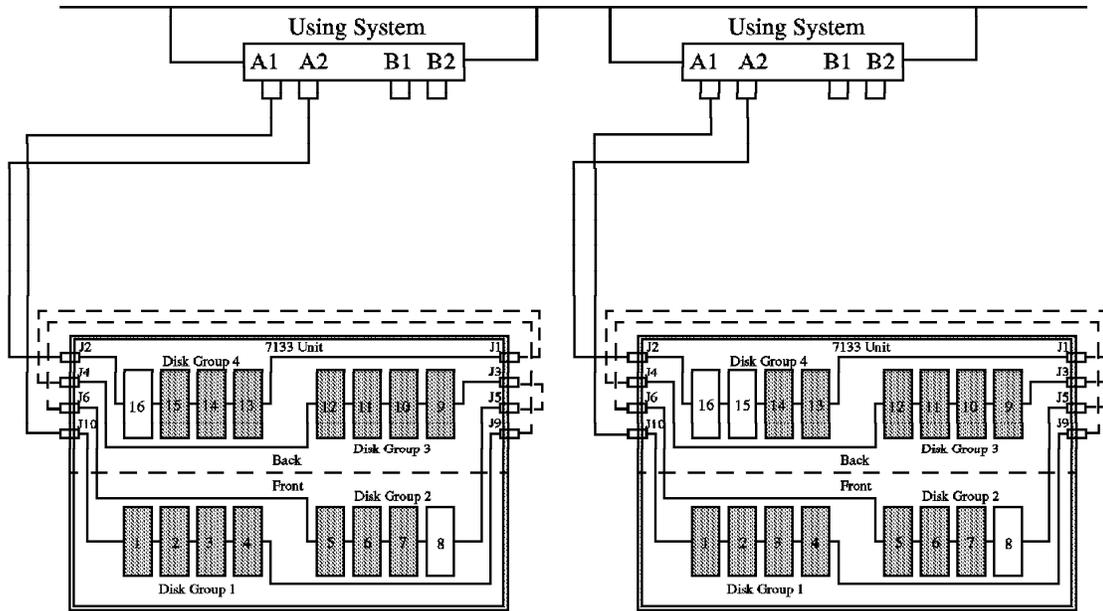


Figure 35. Sequential Access, Capacity 120 GB. This scenario illustrates a design for access to 120 GB.

In this scenario, the workload composition is the same as in the previous one and will require 27 4.5 GB disks. The 4.5 GB disks have been selected due to the cost implications inherent in the number of subsystems that using smaller disks would require.

Since the applications will be driving the disks at a maximum of 1.5 MB per second due to the volume of the data, the loop bandwidth will support 23 disks off a single adapter. Two adapters will therefore be needed to support the total bandwidth requirement. It is best to distribute the disks evenly between the adapters, so two loops, one of 13 disks and one of 14 disks would be optimal.

4.1.2.2 Scenario 2

In this scenario, analysis of existing I/O information shows that the I/O profile for the applications consists of small block size random reads and writes across fairly large files. I/O consists again of 50 percent reads and 50 percent writes.

Storage Capacity Requirements - 20 GB: This can be configured as shown in Figure 36 on page 58. The application is reading small (4 KB) blocks from reasonably large files spread uniformly across the disks.

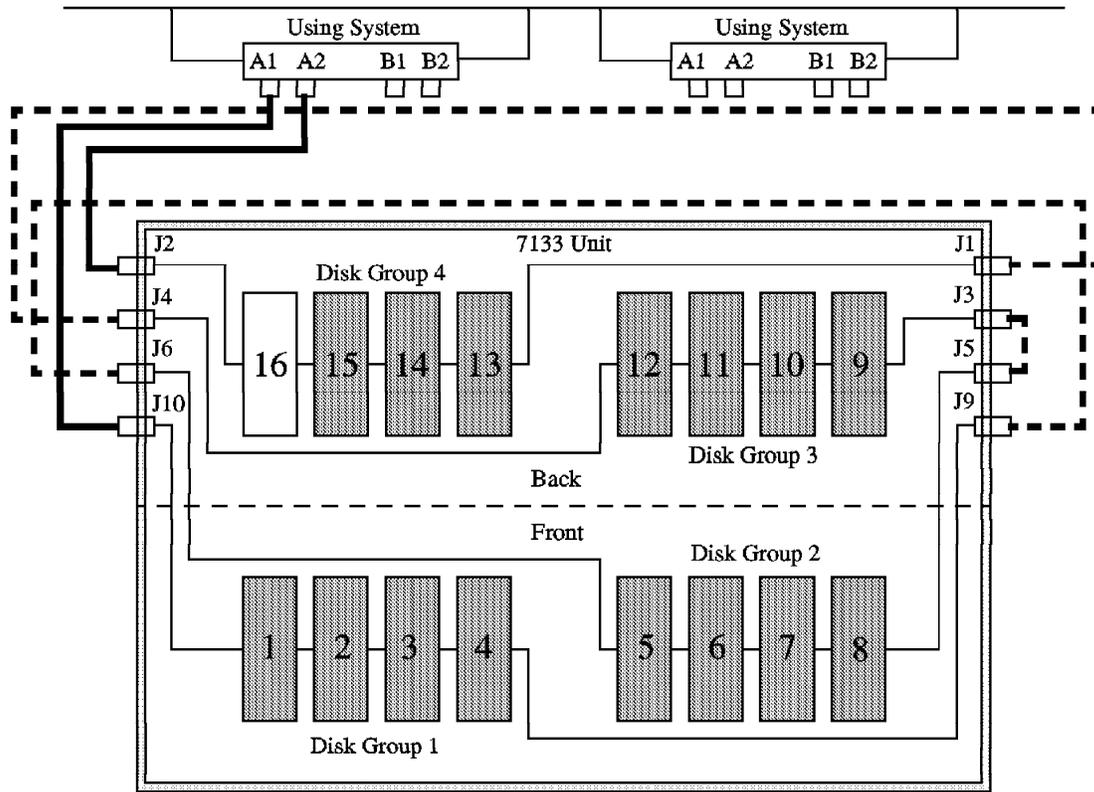


Figure 36. Random Access, Capacity 20 GB. This scenario illustrates a design for access to 20 GB.

With this type of I/O profile, the transaction rate is likely to be of more concern than the throughput. Random I/O involves more disk head seeking and tends to result in more operations per second due to the smaller block sizes involved. For these reasons, the transaction load at the disks and the adapter will be significant.

In this scenario, the applications generate a transaction rate of around 1000 operations per second in total at peak rates. Assuming a maximum rate of around 70 operations per second at a disk for random I/Os, 15 disks would be required. Using 2.2 GB disks will provide a capacity of 33 GB, which is more than is required. It would be possible to arrive at a more exact number by mixing 2.2 and 1.1 GB disks, using 11 x 1.1 GB and 4 x 2.2 GB disks, for example, would give a capacity of 20.9 GB. Since the throughput is much lower (around 4 MB per second), all 15 disks can be placed on the same loop without exceeding the maximum throughput.

Storage Capacity Requirements - 60 GB: This can be configured as shown in Figure 37 on page 59. The operational parameters are the same as for the previous scenario except that the transaction rate is around 1500 operations per second at peak times.

This will require at least 22 disks to support the required transaction rate. In order to fulfill the capacity requirements, this can be implemented with a mix of 2.2 GB and 4.5 GB disks. For example, using 6 x 4.5 GB disks and 16 x 2.2 GB disks will give a capacity of 62.2 GB. Again, the throughput will not approach the bandwidth of the loop (it will be around 6 MB per second) so the disks can all be attached to the same loop, though two subsystems will be required.

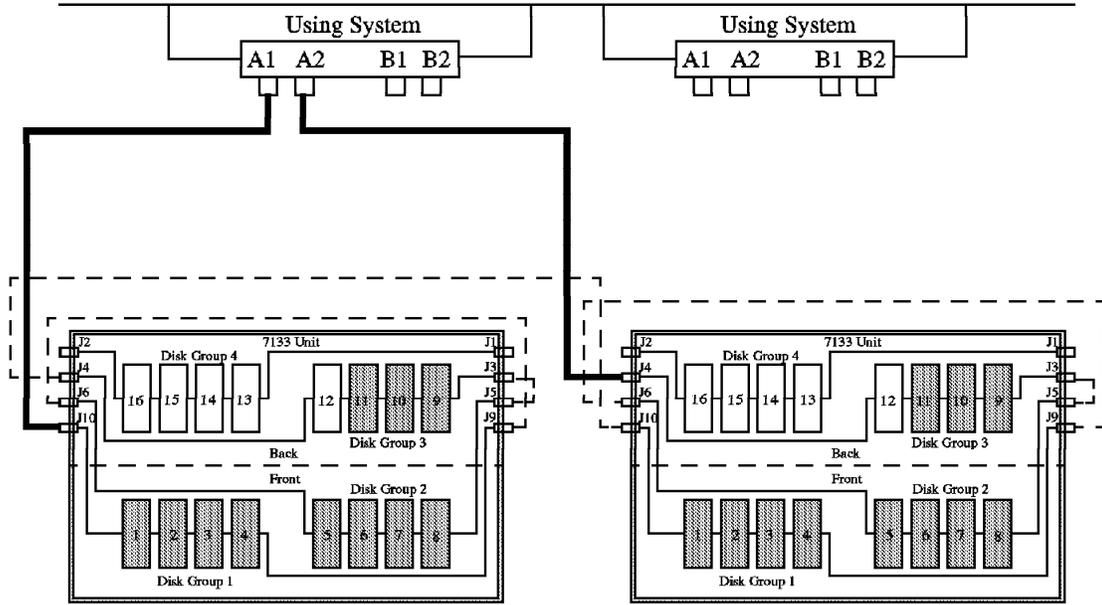


Figure 37. Random Access, Capacity 60 GB. This scenario illustrates a design for access to 60 GB.

Storage Capacity Requirements - 120 GB: This can be configured as shown in Figure 38. In this example, the workload profile is the same as in the previous one with random I/Os in 4 KB blocks to reasonably large files. In this case though, the transaction rate will be around 2000 operations per second.

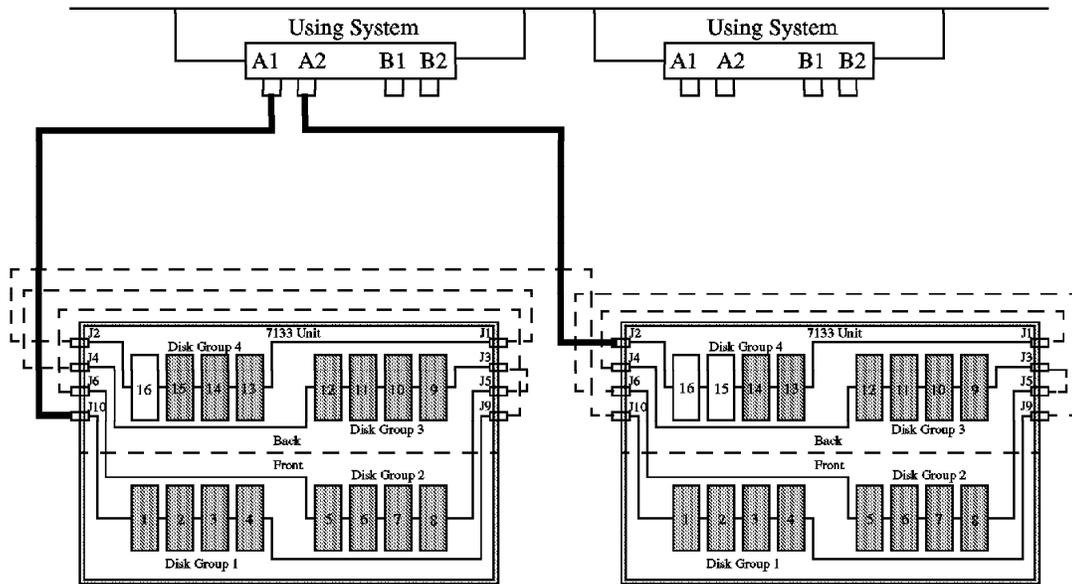


Figure 38. Random Access, Capacity 120 GB. This scenario illustrates a design for access to 120 GB.

In order to achieve the required transaction rate, 29 disks will be required. This can be implemented with 25 x 4.5 GB disks and 4 x 2.2 GB disks. The throughput will be at around 8 MB per second and so should not exceed the adapter limit for a single loop.

4.1.3 Summary

As can be seen, no performance advantage is gained by putting less than four disks on any single loop for sequential I/O profiles, and less than eight for random I/O profiles. More can be configured, depending on the exact nature of the I/O profile, and the total design considerations. Furthermore, the optimum performance is attained by attempting to balance the load across loops where possible. Essentially, the disk bandwidth, loop bandwidth and the number of I/O operations generated per second for the loop and disks must be taken into account when adding disks to a loop.

In situations where the I/O access patterns are not known beforehand, and the I/O profile is mainly sequential using large block sizes, it is best to distribute disks evenly between loops as in scenario 1. Subsequent I/O tracing can reveal any major skew that needs correcting, at which time disks can be moved using the information available in 3.4, "Additional Configurations" on page 42.

For predominantly random I/O, the transaction load should be determined and the loop designed so as not to exceed the transaction rates of the adapters and disks. Depending on the block size and total transaction rate, the throughput of the loop may need to be considered.

There is no direct performance advantage in the size of disk selected though when many files need to be accessed many times, the more actuators that are present, the more concurrency that can be achieved. For this reason, performance may be optimized for the reading and writing of many files by distributing those files across more smaller disks to increase the potential for concurrency.

Location of data on the disks can affect performance though this is an LVM consideration. Locating frequently accessed data near the center of disks generally increases the speed of access to that data.

4.2 Availability Related Scenarios

As discussed earlier, Serial Storage Architecture has availability benefits inherent in its technology. This section will discuss in detail how these benefits can be exploited using a number of scenarios.

4.2.1 Considerations

This section only covers availability in terms of the availability of the subsystem, HACMP is not discussed.

4.2.1.1 Multiple System/Adapter Access

There are a number of configuration possibilities that can be used to implement various levels of high availability.

Single Adapter: The lowest level of availability is achieved with a single SSA adapter in a host system. With this configuration, the native SSA feature of two independent data paths to each disk provides protection against a single cable or peripheral device failure. Should a cable or disk device fail, all other devices on the loop can still be accessed.

Multiple Adapters in a Single Host System: The level of protection can be improved in a configuration with multiple adapters in a single host system. Up to two adapters in a single host system can be on the same loop, and in an environment of this kind, the second adapter can provide another point of redundancy; should an adapter fail, the second adapter will automatically be used to route data to and from the devices on the loop.

Single Adapters in Multiple Host Systems: The next level of availability is achieved with a configuration using single adapters in multiple hosts. Using the new SSA adapter (feature code 6216), up to 8 adapters can be on one loop. Therefore, in this environment, up to eight hosts can be simultaneously connected to the same loop. If HACMP is used in conjunction with this hardware configuration, then host system redundancy is introduced. This scenario can now withstand failure of any one cable, any one peripheral device, or any host system or adapter. The actual level of protection will depend on what component has failed. Certain instances of multiple cable or device failure could be guarded against by distributing the hosts systems evenly around the loop and using HACMP to failover to a host system not cut off by the multiple device/cable failure.

Multiple Adapters in Multiple Host Systems: This is the highest level of availability that is attainable. With multiple adapters in multiple systems and using the HACMP software, the benefits of adapter redundancy can be added to the previous configurations capabilities.

4.2.1.2 Mirroring

Mirroring is implemented at the software level (AIX Logical Volume Manager). However, the implications of mirroring within a loop verses across loops does need to be considered before setting up volume groups and logical volumes. Generally, the best availability is obtained when mirroring across loops; this will be illustrated when the various scenarios are explored. Note that mirroring is assumed to be from disk to disk; there can be no availability benefits mirroring across a disk.

4.2.1.3 Multiple Loops

From an availability viewpoint, to increase availability, it is wise to consider that for n disks, the ideal solution would be to have $(n + 1)$ divided by 2 as the number of loops (rounded down to the nearest whole number). However, this may not always be practical. In some cases, particularly when the number of disks required is small, splitting the disks across loops may not increase availability significantly, unless multiple adapters are also used. In these cases, the cost and complexity must be weighed against the level of availability required. A single loop has only one single point of failure, the adapter. Splitting disks across two loops on the same adapter does not remove this single point of failure.

4.2.1.4 Data Location

Although the placement of the data is controlled at the software level through the AIX operating system LVM, the location on the disks, and which disks are actually to be used for particular data, needs to be decided when the volume group and logical volume arrangement is designed. This is closely linked to mirroring (see 4.2.1.2, "Mirroring").

Generally speaking, the greatest availability will be gained where the data is located on disks in different loops. Thus, a volume group should have its disks spread evenly across the loops available.

4.2.1.5 Multiple Disk Types

As mentioned earlier, SSA allows loops to be created with any of the currently available SSA disks (see 2.1.2.4, “7133 Hardware Schematics and Connection Rules” on page 24). From an availability viewpoint, there are three things to consider:

1. Number of disks

The fewer the number of disks, the less the chances of a subsystem failure. This is simple probability mathematics:

Consider a situation where 20 GB of disk space is required; this could be fulfilled in a number of ways. Two possible solutions would be to use 10 x 2.2 GB disks or 5 x 4.5 GB disks.

With the first situation, there is an associated probability of failure between the disks (A), and the cabling (B), so the total chance of failure can be represented by the formula $(10 \times A) + (4 \times B)$.

With the second solution, the associated probability is $(4 \times A) + (2 \times B)$.

Thus, the first configuration has a higher relative probability of having a problem than the second. So the second configuration has a higher probability of remaining available than the first.

2. Size of disks

Again, using the example of 4.5 and 2.2 GB disks, if one 4.5 GB disk should fail, then the amount of data that will be inaccessible (assuming no mirroring) would be greater than if one 2.2 GB disk should fail.

3. Quorum

The maintenance of disk quorum is a factor in availability; the more disks a volume group consists of, the more chance that the volume group will be able to be varied on if one or more disks are inactive. This is because a volume group requires a majority of the VGDA (Volume Group Descriptor Area) and VGSA (Volume Group Status Area) to be available:

- Single disk

For a single disk volume group, there are two copies of the VGDA and VGSA on the disk. In this instance, if the disk is not available, then obviously the volume group cannot be varied on.

- Two disks

In the case of a two disk volume group, one disk contains two copies of the VGDA/VGSA, and the other has one copy. So if the disk with two copies is unavailable, then the volume group cannot be varied on. If the other disk is not available, then the volume group can still be made available. Thus, there is a 50% chance that the volume group cannot be varied on if one of the disks has a problem.

- Three or more disks

In the case of three (or more) disks, one copy of the VGDA and VGSA will be on each disk. In this situation, any one disk can fail, and quorum can be maintained. Naturally, with more disks in the volume group, the greater the number of disks that can have a problem, and the volume

group still be varied on. In fact, as long as 51 percent of VGDA's are available, a volume group can be varied on.

Bear in mind that so far only volume group availability has been discussed. The data on any failed disks in a volume group that can be varied on will still be unavailable, even though the volume group is active.

Also note that although it is possible to have non-quorum volume groups, it is not recommended to do so as data integrity cannot be assured. If such volume groups are configured, mirroring is advised. See InfoExplorer documentation for more information on quorums.

Taken by themselves, one could use these factors to decide how to set up the subsystem hardware; however, it is rarely the case that one requires only performance or availability from the storage subsystem, usually a combination of both is necessary.

4.2.2 Scenarios

Here are some scenarios designed to illustrate the availability criteria that have been discussed:

4.2.2.1 Scenario 1 - One Host with One Adapter

For this scenario, it must be noted that the one host and adapter card are single points of failure. Therefore, it is necessary to concentrate on maximizing the availability of the subsystem and the data on the disks inside.

Storage Requirement for 20 GB: First, the number of disks required must be decided. This can be as few as 5 x 4.5 GB or as many as 18 x 1.1 GB. If many file systems will be created, one might be tempted to have many smaller disks, so that if a disk is unavailable, the data on an unaffected disk may still be available either as a result of maintaining entire file systems on single disks and/or being able to maintain disk quorum.

However, there are only 16 slots in a single subsystem and a maximum of two loops per adapter. Thus the recommendation is 4 x 4.5 GB disks and 2 x 1.1GB disks, with each loop consisting of 2 x 4.5 and one 1.1 GB disk. Mirroring between the loops should be implemented. If file systems will be relatively small, then it is worth considering implementing each loop with nine 2.2 GB disks, thereby reducing the already tiny chance of a double disk failure, preventing access to data.

With this scenario, there are four paths to access any required data. See Figure 39 on page 64.

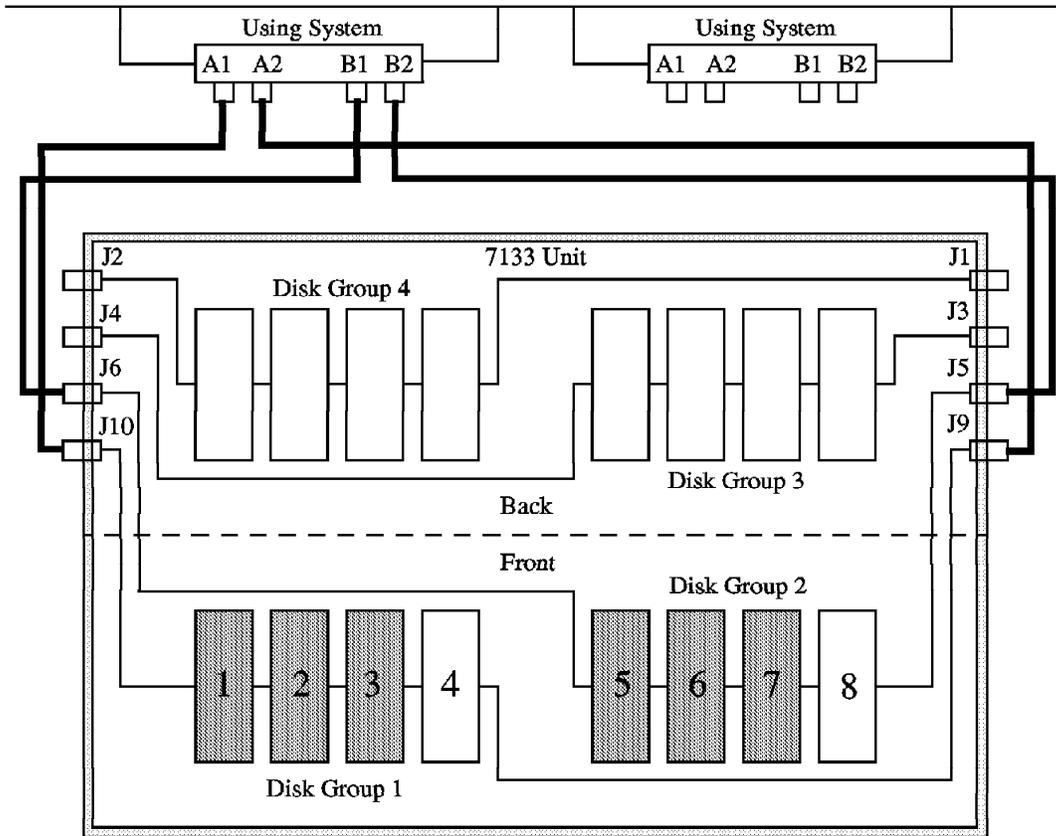


Figure 39. One Host, One Adapter, Capacity 20 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, one adapter and 20 GB of storage.

Note

This scenario is contrary to the recommendation that six disks should go in slots 1 to 6. The first six slots recommendation is from an administrative perspective, not from an availability one. The other scenarios may also not conform to the administrative rules.

Storage Requirement for 60 GB: Bearing in mind that only two loops can be built, to obtain 60 GB of storage will require at least 14 x 4.5 GB disks. These should be set up in such a way that there are seven on each loop, with mirroring implemented across the loops as in the first example.

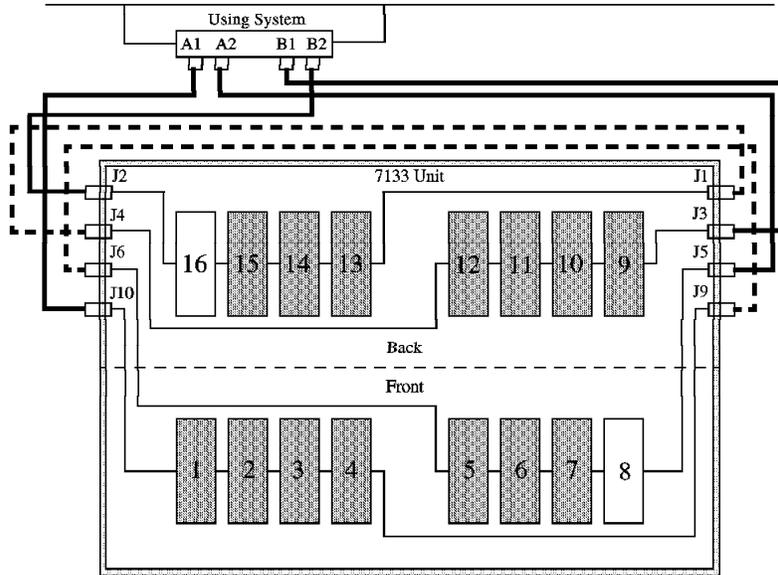


Figure 40. One Host, One Adapter, Capacity 60 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, one adapter and 60 GB of storage.

Storage Requirement for 120 GB: In this scenario, at least two subsystems are required, but the implementation concept is exactly the same (see Figure 41). Again, the design is aimed at having an even number of disks, half on each loop, each loop in a separate subsystem and mirroring implemented between loops.

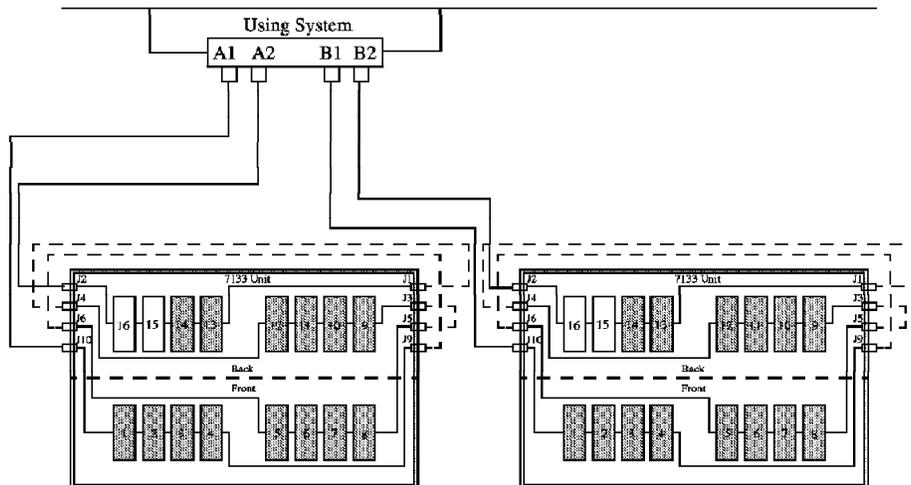


Figure 41. One Host, One Adapter, Capacity 120 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, one adapter and 120 GB of storage.

Note that the connections for the first subsystem run as follows:

- A1 to J10
- J9 to J6
- J5 to J3
- J4 to J1
- J2 to A2

The second subsystem is connected in the same manner to the B ports.

4.2.2.2 Scenario 2 - One Host with More than One Adapter

With the introduction of additional adapter(s), the adapter itself is removed as a single point of failure. This also means that more loops can be used to configure the disks.

Storage Requirement for 20 GB: With two adapters, the opportunity for adapter redundancy now becomes available. With more than two, a very high data/disk availability configuration can be implemented.

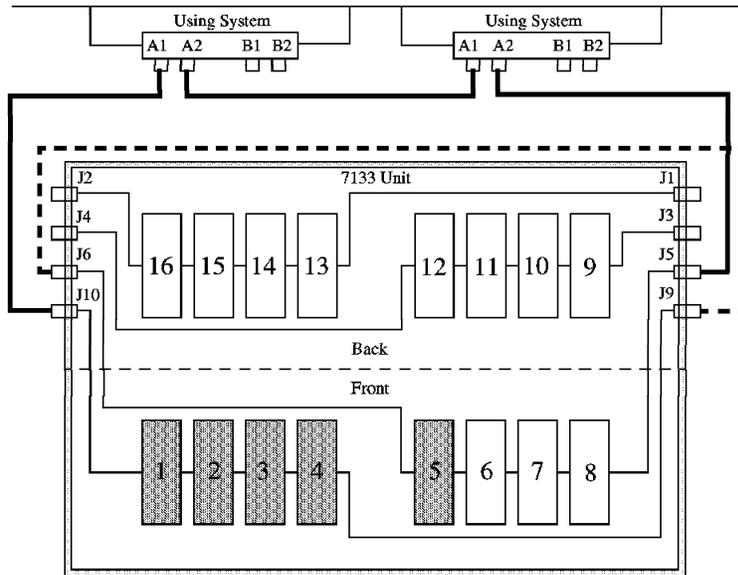


Figure 42. One Host, Two Adapters, Capacity 20 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, two adapters and 20 GB of storage.

This is the standard adapter redundancy configuration.

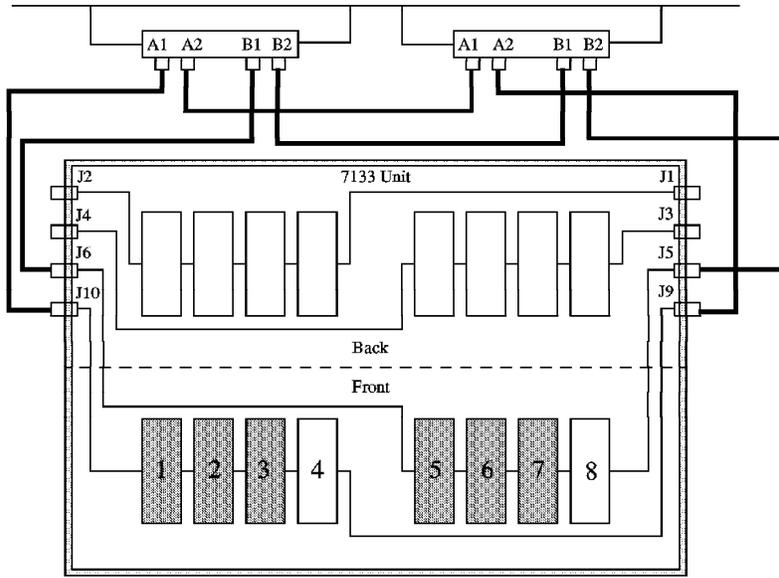


Figure 43. One Host, Two Adapters, Capacity 20 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, two adapters and 20 GB of storage.

This is a higher data availability scenario.

As can be seen from Figure 43, the layout looks similar to the one for a single adapter, except that both loops connect to both adapters, and the adapters communicate directly to each other. So if adapter A cannot talk to disk 5 directly, for whatever reason, then it will communicate via adapter B. This re-routing will occur at the adapter level. This scenario is especially effective where one requires high availability with mirroring. Note also that to even the number of disks per loop for mirroring purposes, the combination of disks here would be 4 x 4.5 GB plus 2 x 2.2 GB. With this alternate scenario, redundancy is implemented at the LVM level by the additional mirrored copy, and is in addition to the adapter redundancy.

With three adapters, the following would be a valid configuration:

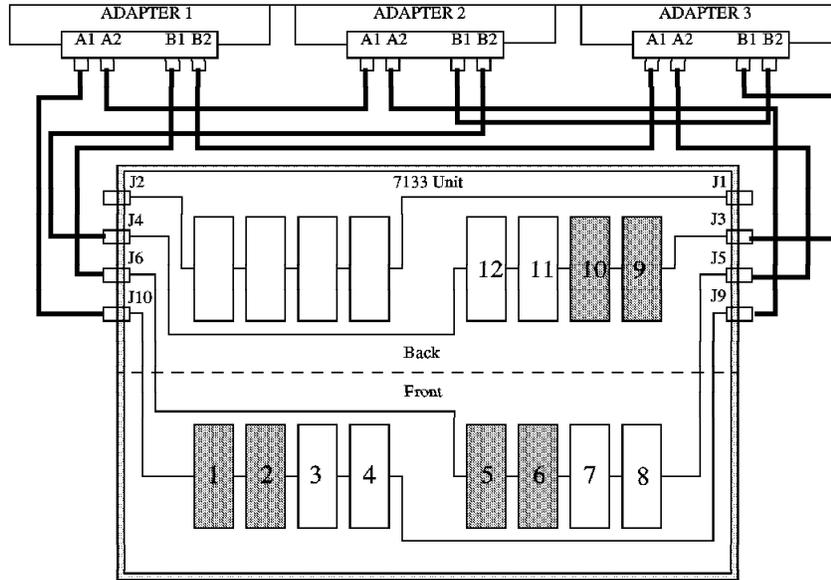


Figure 44. One Host, Three Adapters, Capacity 20 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, three adapters and 20 GB of storage.

This is the highest data availability scenario, where three copies of data are available.

Note that the connections in the diagram are:

- A1 on adapter 1 to J10
- J9 to A2 on adapter 2
- A1 on adapter 2 to A2 on adapter 1
- B1 on adapter 1 to J6
- J5 to A2 on adapter 3
- A1 on adapter 3 to B2 on adapter 1
- B1 on adapter 2 to J3
- J4 to B2 on adapter 2
- B1 on adapter 2 to B2 on adapter 3

Storage Requirement for 60 GB: Again, there are two distinct configurations that will give the same degree of availability in two slightly different ways. These are extensions of the first example (see Figure 45 on page 69 and Figure 46 on page 70).

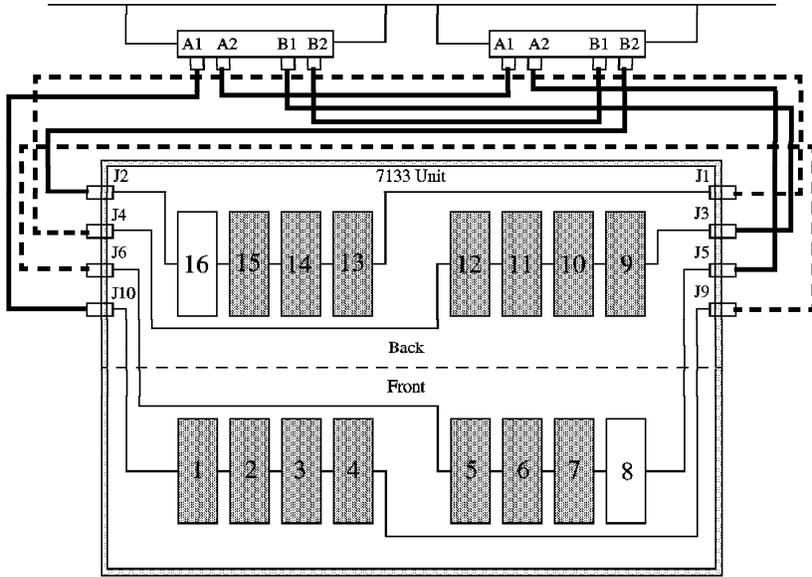


Figure 45. One Host, Two Adapters, Capacity 60 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, two adapters and 60 GB of storage.

This is the standard adapter redundancy scenario.

The connections are:

- A1 on adapter 1 to J10
- J9 to J6
- J5 to A2 on adapter 2
- A1 on adapter 2 to A2 on adapter 1
- B1 on adapter 1 to J3
- J4 to J1
- J2 to B2 on adapter 2
- B1 on adapter 2 to B2 on adapter 1

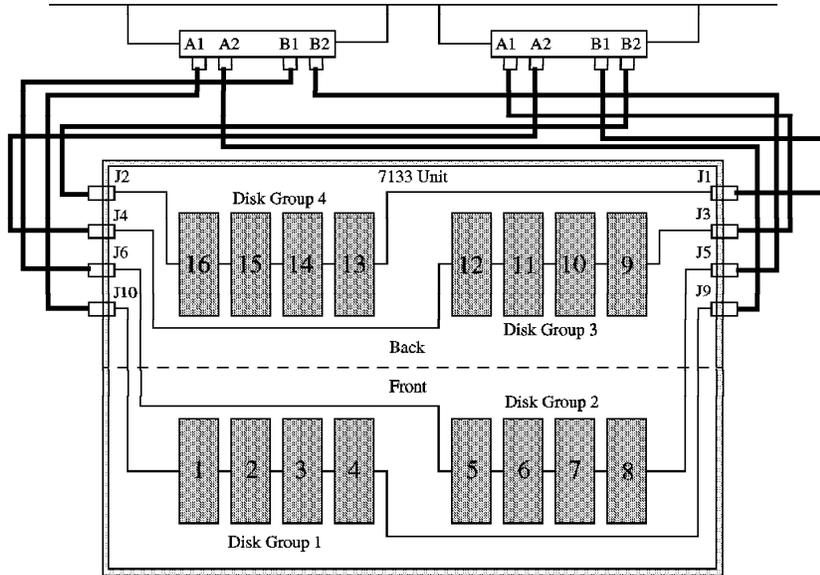


Figure 46. One Host, Two Adapters, Capacity 60 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, two adapters and 60 GB of storage.

Here the connections are:

- A1 on adapter 1 to J10
- J9 to A2 on adapter 1
- B1 on adapter 1 to J6
- J5 to B2 on adapter 1
- A1 on Adapter 2 to J3
- J4 to A2 on adapter 2
- B1 on adapter 1 to J1
- J2 to B2 on adapter 1

In the second diagram, the mirroring should occur across the adapters. The level of availability is the same for both configurations. In the first diagram, 14 x 4.5 GB disks can be used, seven on each loop. In the second diagram, the total number of disks used should be divisible by four, to spread the workload evenly. So the recommended disks to use in this instance would be 12 x 4.5 GB, 2 x 2.2 GB and 2 x 1.1 GB.

Storage Requirement for 120 GB: Again this scenario requires at least two subsystems. The discussion point here is whether to connect one subsystem per adapter, connect both adapters to one subsystem and connect the two subsystems together, or both. From the highest availability viewpoint, the solution must reflect the first scenario, where the disks are split between the subsystems and there are two loops, one for each subsystem set of disks, each loop using both adapters (see Figure 47 on page 71).

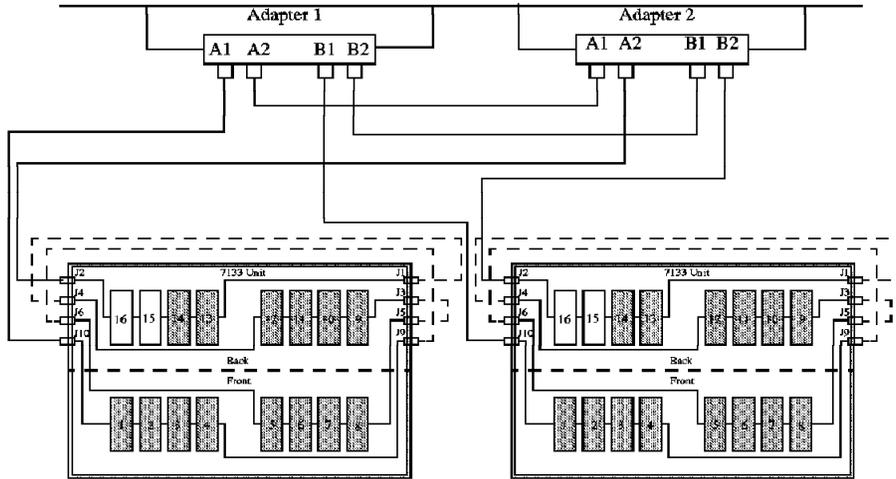


Figure 47. One Host, Two Adapters, Capacity 120 GB. This scenario illustrates a design for high availability of the disk subsystem with one host, two adapters and 120 GB of storage.

Connections for the first subsystem are:

- A1 on adapter 1 to J10
- J9 to J6
- J5 to J3
- J4 to J1
- J2 to A2 on adapter 2
- A1 on adaptor 2 to A2 on adapter 1

The connections for the second subsystem run in exactly the same way for the B ports on the adapters.

4.2.2.3 Scenario 3 - Two Hosts with Two Adapters

As discussed earlier, the previous scenarios had the problem that the host system was a single point of failure. In this example, examples using two host systems and two adapter cards will be considered.

In this scenario there are two loops with five disks each (the size of each disks is 2.2 GB). Each loop runs from one machine/adaptor to the other. Here, any volume groups will consist of an even number of disks, with one in each loop. Mirroring will be across loops. With this scenario, there are four ways of accessing any one disk; the problem will be controlling if and when each of the host systems should be accessing the disks. This scenario would be appropriate for a primary/secondary host situation, where the secondary host acts as a backup for the primary host; should the primary host or adapter fail, it will be relatively easy for the secondary host to assume control.

As with any mirrored system, there is now redundancy in the storage subsystem, and in reality, the storage is only 10 GB (obviously if 20 GB of unique storage space was required, then 4.5 GB disks would have to have been used).

If mirroring is not used, then the single point of failure is the actual disk itself. However, as all of these scenarios are from an availability standpoint, it is assumed that there will be mirroring.

The two hosts/two adapter configurations will look exactly like the ones for the single host/two adapter scenarios, and they typically blend with the HACMP fallback/takeover situation.

4.2.2.4 Scenario 4 - Two Hosts with more than Two Adapters

Obviously, with three adapters in each host, there are benefits to be gained from organizing the disks and the data in such a way that up to two adapters can fail and any disk is accessible. Three mirrored copies would give six different ways in which a piece of data can be accessed. In this case, the host system failing is the point of highest failure, but this is not a single point of failure.

4.2.2.5 Scenario 5 - More than Two Hosts

With the recently announced SSA adapter (feature code 6216), it is now possible to have up to eight adapters on a single loop. From the point of view of availability, this means that very high levels of redundancy can be achieved when used in conjunction with HACMP; effectively, up to eight hosts can be attached to a single loop.

4.2.2.6 Summary

As has been seen from the various scenarios, from an availability standpoint, some of the configurations can be anything but standard (most actually contradict the administrative rules for how disks should be configured). As with most things, an acceptable balance between performance and availability should be sought, and the configuration used should be based on this.

Cost will probably be a large influencing factor in the decision taken. In most of the availability configurations, mirroring was recommended though this may not be an important factor in the cost, performance and availability equation. In the next section, combining performance and availability will be examined, to get the best of both worlds in a cost effective manner.

4.3 Performance and Availability

This section looks at the situations where both performance and availability are important considerations. The criteria that are involved in designing a solution for optimum performance or for optimum availability have already been discussed in the previous sections, so here, the focus will be on any trade-offs that need to be made.

4.3.1 Considerations

The considerations that were discussed in the sections on performance and availability will now be re-examined in the light of designing solutions that provide an optimum level of both requirements.

4.3.1.1 Bandwidth

The main restriction on design that bandwidth places is on the maximum number of physical disks that can be connected to a single loop. The exact number depends upon the I/O profile of the applications that are using the disks. For an application that is generating large block size sequential I/O, the throughput may become the limiting factor. If so, multiple adapters or loops may need to be considered. If the profile is small block size random I/O, then the transaction rate, or number of I/O operations per second will probably be the critical factor.

It will be necessary to ensure that the rate does not exceed the maximum for an adapter and that sufficient disks are available to support the requirement.

1. Cost implications

Restricting the maximum number of disks on a loop may mean an increase in cost, particularly if a large capacity must be mirrored. However, if cost is not an issue, then both performance and availability can be enhanced (depending on the configuration) by increasing the number of adapters and loops.

2. Double failures

As the number of disks on a loop increases, so does the number of disks that will be isolated in the extremely unlikely event of a double disk or cable failure. Considering Figure 41 on page 65, if disks 1 and 14 should fail simultaneously, then there will no longer be a path to the 12 disks in between. Note that the chances of this happening are remote in the extreme, but it may be an availability consideration, in which case a policy of reducing the number of disks on a loop for performance reasons will also have availability benefits.

The key point here is that optimum performance is achieved from balancing the load by spreading the disks across loops and adapters so as to avoid exceeding the bandwidth capacity of loops, adapters, CPU and MicroChannel bus. Having multiple adapters on a loop can also improve performance up to a certain point. A good design that satisfies both requirements can therefore be arrived at (see Figure 45 on page 69 for an example).

4.3.1.2 Multiple Adapter Access

The number of adapters on a particular loop can improve the bandwidth available on a loop. With two adapters on the same host with a single MicroChannel for example, the available bandwidth can increase to 45 MB per second. There are also obvious availability improvements as a result of using multiple adapters.

With multiple systems, the number of disks on a loop split between two systems could be doubled without meeting the bandwidth limit of the loop. Half of the disks on the loop could communicate with one system adapter while the other half could communicate with the second.

4.3.1.3 Data Path

There are two elements that govern the data path to a disk:

1. Adapter ownership

By default, ownership of disks in a loop with two adapters, where both adapters are in the same host system, is designed to minimize the path to each disk; this is optimum from a performance point of view. If the adapters are in different host systems then it is up to the user to locate the data such that the path length is minimized.

2. Disk position

The position of a disk on a loop will have no impact on the performance of that disk at low utilization levels. As the utilization of the loop increases towards the maximum, those disks furthest from the adapter that they are communicating with will receive slightly improved service. There is also a small availability consideration: the longer loops are, the more disks that will

be isolated in the unlikely event of a double failure at the ends of the loop. The risk of such an occurrence will have to be traded off against the cost of multiple loops. This may not be significant in the case of a single adapter, only requiring additional cables to use both sets of ports, and is part of the reason for balancing loads across the ports.

4.3.1.4 Data Location

The location of data on a disk has no direct impact on availability. Thus locating data that has high performance requirements at the center of disks will not affect the availability of that data.

4.3.1.5 Multiple Disk Types

The types of disk used may have an impact on performance, depending on the sizes and frequencies of data access. There is only a very small impact on availability. Using smaller disks to increase concurrency will require more physical disks to meet the capacity requirements. This will increase the chances of a failure by a very small amount. Generally speaking, this is not an issue as the probability increase is so small. Using mirroring provides protection against the chances of double failure.

4.3.1.6 Multiple System Access

The number of systems attached to a loop will directly affect availability. A single system is a single point of failure, so having two adapters in a loop, each in a separate system, provides a higher level of availability.

The number of systems has no direct effect on each individual systems performance though, although the performance of the loop itself would be greater; obviously, data on a specific disk will usually only be able to be accessed from one system at a time.

4.3.1.7 Mirroring

Availability is greatly improved through the use of mirroring. In terms of performance, however, the only impact is in terms of how the mirroring is actually implemented. Mirroring can be parallel or sequential:

1. Parallel

In parallel mirroring, both mirror copies are updated simultaneously which provides the best performance but the lowest level of protection in the event of a failure during mirror update.

2. Sequential

In sequential mirroring, each copy is not updated until the previous copy update has completed. This ensures consistency and provides the highest levels of availability, but the worst performance.

Both parallel and sequential mirroring only affect the write performance. Generally speaking, read performance improves with mirroring as the request will be sent to the disk that can provide the fastest response for the data required.

Both types of mirroring will increase the I/O bandwidth and CPU utilization for writes, and this should be taken into consideration as it could be significant in some environments. Sequential mirroring will provide slightly higher protection, but with a small extra performance cost. This is the trade-off.

4.3.1.8 Multiple Loops

The number of loops in a configuration has a direct impact on availability. More loops implies greater redundancy, which reduces the chances of a failure preventing access to any disks. This approach should be used in conjunction with mirroring to further reduce the chances of losing access to any data. From a performance standpoint, if the number of disks required would exceed the bandwidth limit for a loop, then having multiple loops will be necessary to maintain optimum performance.

There is no real trade-off required here as multiple loops will not impact performance, if required for availability reasons, as long as the loops can be in the same host system.

4.3.2 Summary

Thus, the major points that need to be considered when designing for performance and availability are:

1. Bandwidth

Ensure that the number of disks on a loop does not exceed the bandwidth of the loop.

2. Transactions

Ensure that the number of transactions per second generated by the applications using a loop does not exceed the overall maximum that the adapter is capable of supporting. In addition, if any disks are operating at the maximum number of I/O operations per second that they support, the changes are that the performance of the loop will not be optimal and utilizing more disks (actuators) would improve performance.

3. Multiple adapters

Two adapters will improve availability of a loop through redundancy. The bandwidth available will also be higher.

4. Mirroring

Ensure that the mirroring design provides the necessary levels of availability and performance. This will be a trade-off.

Once these points have been taken into consideration in the design process, none of the other factors should affect availability or performance negatively.

Chapter 5. Problem Determination and Resolution

This chapter looks at the tools available for problem determination and resolution. This will include looking at the SSA service aids provided and describing the uses of the various components. The second section provides examples of various potential problems and their resolution.

5.1 SSA Service Aids

This section looks at the SSA service aids provided to assist in servicing the 7133 subsystem. The service aids are accessed from the standard AIX diagnostics menu and provide the following functions:

1. Set Service Mode

This option allows the location of a specific SSA disk drive within a loop to be determined and the drive to be removed from the configuration, if required.

2. Link Verification

This option allows the operational status of a link to be determined.

3. Configuration Verification

This option allows the relationships between physical (pdisks) and logical (hdisk) disks to be displayed.

4. Format Disk

This option allows SSA disk drives to be formatted.

5. Certify Disk

This option provides the capability to test whether data on an SSA disk drive can be read correctly.

6. Identify

This facility is accessible from any of the previous options and provides the capability to determine the physical location of SSA disk drives within a subsystem.

The rest of this section will look at these functions in more detail.

5.1.1 Starting the SSA Service Aids

The SSA service aids are accessed from the standard AIX diagnostics menu. This menu is invoked with the `diag` command.

Important

The service aids should not be run on more than one system connected to the same loop at any one time.

Invoking the diagnostics produces the following menu:

FUNCTION SELECTION

801002

Move cursor to selection, then press Enter.

Diagnostic Routines

This selection will test the machine hardware and detect any hardware problems. A problem will be indicated by a SRN (Service Request Number). The SRN will allow a service representative to quickly determine what parts are required to repair the machine.

Service Aids

This selection will look at the machine configuration, exercise external interfaces, format media, look at past diagnostic results, control what resources are tested, check out media, etc.

Advanced Diagnostic Routines

This selection will normally be used by the service representative.

System Exerciser

This selection will test resources running in an overlap mode.

F3=Cancel

F10=Exit

Select **Service Aids** which will produce the following screen:

SERVICE AIDS SELECTION

802001

Move cursor to selection, then press Enter.

[TOP]

Service Hints

This selection displays service hints and errata information about the maintenance package.

Display Previous Diagnostic Results

This selection displays the results of previous diagnostic runs.

Display or Change Configuration or Vital Product Data (VPD)

This selection displays or changes the Configuration or VPD.

Display or Change Diagnostic Test List

This selection displays the resources tested by Diagnostics and allows the resources to be added or deleted from the Diagnostic Test List.

Disk Media

This selection provides the format disk and certify disk service aids.

Diskette Media

This selection provides a tool for checking out a diskette.

Local Area Network Service Aid

This service aid is used to exercise the Local Area Network

[MORE...39]

F3=Cancel

F10=Exit

Scroll forward, and select **SSA Service Aids** from this list of service aids. This will produce the following screen:

```
SSA SERVICE AIDS 802380

Move cursor onto selection, then press <Enter>.

Set Service Mode
Link Verification
Configuration Verification
Format Disk
Certify Disk

F3=Cancel      F10=Exit
```

The SSA Service Aids are now ready to be used.

5.1.2 Using the Service Aids

Each service aid will now be examined in turn and its features and functions described.

5.1.2.1 Set Service Mode

The Set Service Mode option allows the specific location of a particular disk drive within a subsystem to be physically located and then removed from the configuration, if required. Before a disk can be removed from a subsystem, it must be placed in service mode. This prevents any further data traffic to the disk, checks that the SSA loop is not already broken elsewhere, and causes the subsystem check light of the subsystem containing the disk and the check light on the disk itself to light for identification purposes.

Note
The disk must also be made unavailable to any system that is using it prior to this operation. Ensure that any file systems mounted and using this disk are unmounted prior to placing it in service mode.

The procedure can take place while normal operations continue to the rest of the disks in the SSA subsystem. Once the check light on the disk drive is lit, the disk can be removed from the subsystem with no effect on any other components in the subsystem.

Important
If the disk is to be removed for any length of time, it should be replaced with either another drive or a dummy disk drive module to maintain the correct cooling air flow within the subsystem.

In order to perform this process, use the following steps:

1. Select **Set Service Mode** from the SSA Service AIDS menu.

This will produce a list of the physical drives known to this system:

```

SET SERVICE MODE 802381

Move cursor onto selection, then press <Enter>.

pdisk1    5AEA1061  00-03-P  4GB SSA C Physical Disk Drive
pdisk2    5AEA10C8  00-03-P  2GB SSA C Physical Disk Drive

F3=Cancel      F10=Exit

```

The columns displayed have the following meaning:

Column 1 The resource name of the physical disk.

Column 2 The SSA serial number allocated to the physical disk drive.

Column 3 The AIX location code of the physical disk drive AB-CD-P, where:

A Always zero.

B Always zero.

C The system I/O bus identifier.

D The SSA adapter position on the I/O bus (the slot number).

P Indicates and SSA physical disk drive.

Column 4 The description of the physical disk drive.

2. Select the drive that is to be placed in service mode.

Using the cursor arrow keys, highlight the physical drive that is to be put in service mode, and then press the **Enter** key. This will produce the following screen:

```

SET SERVICE MODE 802382

pdisk1    5AEA1061  00-03-P  4GB SSA C Physical Disk Drive

Move cursor onto selection, then press <Enter>.

+ Set or Reset Identify.
  Select this option to set or reset the Identify indicator
  on the disk drive.
> Set or Reset Service Mode.
  Select this option to set or reset Service Mode on the
  disk drive.
  ENSURE THAT NO OTHER HOST SYSTEM IS USING THIS DISK DRIVE
  BEFORE SELECTING THIS OPTION.

F3=Cancel      F10=Exit

```

Two options are shown:

a. Set or Reset Identify

Choose this option to toggle the identify function. When enabled, this option causes the check light on the selected disk to flash. The following screen is shown when this option is selected:

```
SET SERVICE MODE 802381

Move cursor onto selection, then press <Enter>.

+ pdisk1 5AEA1061 00-03-P 4GB SSA C Physical Disk Drive
  pdisk2 5AEA10C8 00-03-P 2GB SSA C Physical Disk Drive

F3=Cancel F10=Exit
```

The + beside the disk indicates that it is in identify mode. Note that it is possible to use the cursor arrow keys here to toggle identify mode for any of the other disks listed. Highlight the required disk and press the enter key. Again the presence of the + symbol beside a disk shows that it is in identify mode and that its check light will be flashing.

b. Set or Reset Service Mode

Choose this option to toggle service mode. When enabled, the selected disk is placed into service mode and may be physically removed from the subsystem. Service mode is indicated by the check light being lit on the disk.

Important

Note that when this option is selected, communication with the disk is disabled. For this reason, ensure that no other system is using the drive before placing it in service mode.

Selecting this option will produce the following screen:

```
SET SERVICE MODE 802381

Move cursor onto selection, then press <Enter>.

> pdisk1 5AEA1061 00-03-P 4GB SSA C Physical Disk Drive
  pdisk2 5AEA10C8 00-03-P 2GB SSA C Physical Disk Drive

F3=Cancel F10=Exit
```

The > beside the disk indicates that it is in service mode. Note that it is possible to use the cursor arrow keys here to toggle service mode for any of the other disks listed. Highlight the required disk, and press the

Enter key. Again, the presence of the > symbol beside a disk shows that it is in service mode and that its check light will now be on.

Returning to the set service mode menu will indicate the new status of the disks:

```
SET SERVICE MODE                                     802381

Move cursor onto selection, then press <Enter>.

> pdisk1      5AEA1061  00-03-P  4GB SSA C Physical Disk Drive
+ pdisk2      5AEA10C8  00-03-P  2GB SSA C Physical Disk Drive

F3=Cancel          F10=Exit
```

In this example, disk pdisk1 is in service mode, and disk pdisk2 is in indicate mode. Exiting from the set service mode screen will reset all disks to normal mode, and a popup message to this effect will be displayed on leaving.

3. Repeatedly press the **F3** key to return to the SSA Service AIDS menu.

Note

1. Only one disk at a time can be in service mode. Once one disk has been placed in service mode, attempts to place another in service mode will fail.
2. If the drive selected is not in a closed SSA loop or at the end of a string (a valid configuration), then the operation will fail, and an error message will be displayed. In this case, use the Link Verification option to identify any problems (see 5.1.2.2, "Link Verification").
3. If any file systems are currently mounted on the disk selected, then the operation will also fail. Use the Configuration Verification option to identify the hdisk that requires file systems to be unmounted (see 5.1.2.3, "Configuration Verification" on page 85).

5.1.2.2 Link Verification

This option allows various status information about a link and the physical devices on the link to be determined. The following information can be displayed:

1. SSA link status
Breaks in an SSA link can be located.
2. SSA disk status
The operational status of SSA drives can be displayed.
3. SSA power status
Any problems with power and cooling modules can be displayed.

In order to display this information, use the following steps:

1. Select **Link Verification** from the SSA Service AIDS menu.

This will produce a screen listing all of the SSA adapters attached to this system:

```
LINK VERIFICATION 802385

Move cursor onto selection, then press <Enter>.

  ssa0      00-03  SSA Adapter

F3=Cancel      F10=Exit
```

In this case, there is only one adapter installed. The fields displayed have the following meaning:

Column 1 The system resource name for the adapter.

Column 2 The AIX location code of the SSA adapter AB-CD, where:

- A* Always zero
- B* Always zero
- C* The system I/O bus identifier
- D* The SSA adapter position on the I/O bus (the slot number)

Column 3 The description of the adapter.

2. Select the adapter to display link information for.

Using the cursor arrow keys, highlight the adapter that is on the link that needs verifying, and press the **Enter** key. This will produce the following screen:

```
LINK VERIFICATION 802386

SSA Link Verification for:
  ssa0      00-03  SSA Adapter

To Set or Reset Identify, move cursor onto selection, then press <Enter>

  Physical      Serial#      Adapter Port      Status
                A1  A2  B1  B2
  pdisk2        5AEA10C8    0  1              Good
  pdisk1        5AEA1061    1  0              Good

F3=Cancel      F10=Exit
```

The columns displayed have the following meanings:

Physical The resource name of the physical disk. Note that this field may also contain the following values:

????? This indicates a point at which the SSA loop has broken and no further information can be determined about devices beyond it on this port.

***** This indicates an unconfigured device which may be any one of the following:

a. Another SSA adapter on this loop

The adapter may be in this or another system.

b. An unknown SSA device

The unknown device may be something physically unknown to the software or an as yet unconfigured SSA disk drive. The configuration manager needs to be run (cfgmgr) in order to configure the new drive.

Serial# The SSA serial number allocated to the physical disk drive.

Adapter Port The adapter ports to which the disks are connected. The adapter has four ports, A1, A2, B1 and B2, and the disks will be in loops on the A or B ports. The figures in the columns under each port designator indicate the relative position of the disk on the loop. Thus it can be seen that pdisk2 is in position zero (the first disk) on port A1 and in position one (the second disk) on port A2.

Status This indicates the current status of the disk and can be one of the following four values:

Good The disk drive is functioning correctly.

Failed The disk drive has failed. Use the procedures described in 5.2.2, "Using the Service Aids to Resolve Problems" on page 94 to resolve this problem.

Power The disk drive has detected a failure in the redundant power supply.

Reserved The disk drive is in use on another system.

3. Select the drive for identify.

Use the cursor arrow keys to select the drive that needs to be physically located, and press the **Enter** key. This toggles the identify function described in 5.1.2.1, "Set Service Mode" on page 79. When enabled, a + symbol will be displayed next to the drive.

Exiting the link verification function will terminate any identify operations that are running.

4. Repeatedly press the **F3** key to return to the SSA Service AIDS.

Note

Any changes made to the configuration may take up to 30 seconds to be recognized by the software. Furthermore, to allow changes to be reflected on the Link Verification screen, exit from this option, and then reselect it.

5.1.2.3 Configuration Verification

This option enables the relationships between physical (pdisks) and logical (hdisks) disks to be determined. Along with this information, the physical connection information for the disk drives is also displayed. This includes the port numbers that the disk is connected to and its position in the loop from each port.

In order to display this information, use the following steps:

1. Select **Configuration Verification** from the SSA Service AIDS menu.

This will produce a list of all of the pdisks and hdisks known to this system:

```
CONFIGURATION VERIFICATION                                802390

Move cursor onto selection, then press <Enter>.

hdisk1    5AEA1061  00-03-L  SSA Logical Disk Drive
pdisk1    5AEA1061  00-03-P  4GB SSA C Physical Disk Drive
pdisk2    5AEA10C8  00-03-P  2GB SSA C Physical Disk Drive
hdisk2    5AEA10C8  00-03-L  SSA Logical Disk Drive

F3=Cancel          F10=Exit
```

The columns displayed have the following meaning:

Column 1 The resource name of the physical disk or logical disk.

Column 2 The SSA serial number allocated to the physical disk drive.

Column 3 The AIX location code of the disk drive AB-CD-E, where:

- A* Always zero.
- B* Always zero.
- C* The system I/O bus identifier.
- D* The SSA adapter position on the I/O bus (the slot number).
- E* A P indicates an SSA physical disk drive and an L indicates a logical disk drive.

2. Select the entity for which information is required.

Use the cursor arrow keys to highlight the pdisk or hdisk about which information is required, and then press the **Enter** key. The information displayed depends on whether an hdisk or a pdisk is selected:

- a. Selecting an hdisk

On pressing the enter key, the following information is displayed for an hdisk:

```

CONFIGURATION VERIFICATION 802391

  hdisk2      5AEA10C8  00-03-L  SSA Logical Disk Drive

To set or reset Identify, move cursor onto selection, then press <Enter>

Physical      Serial#      Adapter      Port      SSA_Addr      Status
-----
pdisk2        5AEA10C8      00-03      A1         0             Good
              5AEA10C8      00-03      A2         1             Good

F3=Cancel          F10=Exit

```

As can be seen, the pdisk associated with this hdisk is displayed along with several columns of further information:

- Physical* The resource name of the physical disk.
- Serial#* The SSA serial number allocated to the physical disk drive.
- Adapter* The AIX location code of the SSA adapter AB-CD, where:
 - A* Always zero
 - B* Always zero
 - C* The system I/O bus identifier
 - D* The SSA adapter position on the I/O bus (the slot number)
- Port* The adapter ports to which the disks are connected. The adapter has four ports, A1, A2, B1, and B2, and the disks will be in loops on the A or B ports.
- SSA_Addr* The relative position of the disk from this port. In this example, it can be seen that pdisk2 is in position zero (the first disk) on port A1 and in position one (the second disk) on port A2.
- Status* This indicates the current status of the disk and can be one of the following four values:
 - Good** The disk drive is functioning correctly.
 - Failed** The disk drive has failed. Use the procedures described in 5.2.2, "Using the Service Aids to Resolve Problems" on page 94 to resolve this problem.
 - Power** The disk drive has detected a failure in the redundant power supply.
 - Reserved** The disk drive is in use on another system.

The identify function can be invoked for this physical disk from this screen. While the disk is highlighted, pressing the **Enter** key will toggle the identify function described in 5.1.2.1, "Set Service Mode" on page 79. When enabled, a + symbol will be displayed next to the drive.

Exiting the configuration verification function will terminate any outstanding identify operations for physical disks.

b. Selecting a pdisk

On pressing the Enter key, the following information is displayed for a pdisk:

```
CONFIGURATION VERIFICATION                                     802392
pdisk2      5AEA10C8    00-03-P    2GB SSA C Physical Disk Drive
Move cursor onto selection, then press <Enter>.
hdisk2      5AEA10C8    00-03-L    SSA Logical Disk Drive      Good
F3=Cancel          F10=Exit
```

As can be seen, the hdisk associated with the pdisk is displayed along with several columns of information:

Column 1 The resource name of the logical disk.

Column 2 The SSA serial number allocated to the physical disk drive.

Column 3 The AIX location code of the logical disk drive AB-CD-L, where:

A Always zero.

B Always zero.

C The system I/O bus identifier.

D The SSA adapter position on the I/O bus (the slot number).

L Indicates an SSA logical disk drive.

Selecting the associated hdisk (by pressing the **Enter** key) will have the same result as selecting an hdisk on the previous screen. This has already been discussed.

3. Repeatedly press the **F3** key until the SSA Service AIDS menu is displayed again.

5.1.2.4 Format Disk

This option formats an SSA disk. The format procedure reinitializes the disk.

Important

This procedure will erase any data on the disk. Ensure that any important information on the disk is backed up prior to using this procedure.

In order to format an SSA disk, use the following steps:

1. Select **Format Disk** from the SSA Service AIDS menu.

This will produce a screen listing all of the physical disks owned by this system:

```
FORMAT DISK 802395

Move cursor onto selection, then press <Enter>.

pdisk1 5AEA1061 00-03-P 4GB SSA C Physical Disk Drive
pdisk2 5AEA10C8 00-03-P 2GB SSA C Physical Disk Drive

F3=Cancel F10=Exit
```

The columns displayed have the following meaning:

Column 1 The resource name of the physical disk.

Column 2 The SSA serial number allocated to the physical disk drive.

Column 3 The AIX location code of the disk drive AB-CD-E, where:

A Always zero.

B Always zero.

C The system I/O bus identifier.

D The SSA adapter position on the I/O bus (the slot number).

P Indicates an SSA physical disk drive.

2. Select the drive to be formatted.

Using the cursor arrow keys, highlight the drive that is to be formatted, and press the **Enter** key. This will produce the following screen:

```
FORMAT DISK 802396

pdisk1 5AEA1061 00-03-P 4GB SSA C Physical Disk Drive

Move cursor onto selection, then press <Enter>.

+ Set or Reset Identify Mode.
  Select this option to set or reset the Identify indicator
  on the disk drive.
  Format.
  Select this option only if you are sure that you have selected
  the correct disk drive.
  FORMATTING DESTROYS ALL DATA ON THE DISK DRIVE.

F3=Cancel F10=Exit
```

Details on the selected disk are displayed and two options provided:

a. Set or Reset Identify Mode

As described previously in 5.1.2.1, "Set Service Mode" on page 79, this option allows the disk to be physically located by flashing the check light

on the disk and lighting the check light on the 7133 subsystem containing the disk. Use this option to ensure the correct disk will be formatted.

Note that this option toggles identify mode, but leaving the format service aid will always reset any outstanding identify operations.

b. Format

Select this option to format the selected disk. Format time depends upon the size of the disk:

4.5 GB Approximately 30 minutes

2.2 GB Approximately 16 minutes

1.1 GB Approximately eight minutes

On completion, the disk will be formatted and all data erased. Pressing **Enter** produces the following screen:

```
FORMAT DISK                                     802396
*** WARNING ***
Disk drive selected is:
  pdisk1   5AEA1061  00-03-P  4GB SSA C Physical Disk Drive

All data on the physical disk will be lost.
Note: Before formatting a physical disk, ensure that the disk has been
      powered on for at least 30 minutes.
Do you want to continue with the format operation ?

NO
YES

F3=Cancel          F10=Exit
```

Use the cursor arrow keys to highlight YES to continue the format and NO to abort this procedure. Press **Enter** to initiate the action. Selecting YES will produce the following screen:

```
FORMAT DISK                                     802399
Formatting:-
  pdisk1   5AEA1061  00-03-P  4GB SSA C Physical Disk Drive

                          Format is [  0 ] % Complete
                          [                               ]
Please stand by.

F3=Cancel          F10=Exit
```

The percentage complete value will increment and a status bar grow across the screen to indicate the progress of the operation. On completion, a success message should be displayed. Press the **Enter** key to exit from this screen.

3. Repeatedly press the **Enter** key until the SSA Service AIDS menu is again displayed.

5.1.2.5 Certify Disk

This option will verify that data can be correctly read from the disk selected. In order to initiate this function, use the following steps:

1. Select **Certify Disk** from the SSA Service AIDS menu.

This will produce a list of all of the physical disks owned by the system:

```
CERTIFY DISK                                     802404

Move cursor onto selection, then press <Enter>.

pdisk1      5AEA1061  00-03-P  4GB SSA C Physical Disk Drive
pdisk2      5AEA10C8  00-03-P  2GB SSA C Physical Disk Drive

F3=Cancel          F10=Exit
```

The columns displayed have the following meaning:

Column 1 The resource name of the physical disk.

Column 2 The SSA serial number allocated to the physical disk drive.

Column 3 The AIX location code of the disk drive AB-CD-P, where:

A Always zero.

B Always zero.

C The system I/O bus identifier.

D The SSA adapter position on the I/O bus (the slot number).

P Indicates an SSA physical disk drive.

2. Select the disk that is to be certified.

Using the cursor arrow keys, highlight the disk that is to be certified, and press the **Enter** key. This will produce the following screen:

```

CERTIFY DISK 802405

  pdisk1    5AEA1061  00-03-P  4GB SSA C Physical Disk Drive

Move cursor onto selection, then press <Enter>.

+ Set or Reset Identify Mode.
  Select this option to set or reset the Identify indicator
  on the disk drive.
Certify.
  Select this option to start the Certify operation.

F3=Cancel      F10=Exit

```

Details on the selected disk are displayed and two options provided:

a. Set or Reset Identify Mode

As described previously in 5.1.2.1, “Set Service Mode” on page 79, this option allows the disk to be physically located by flashing the check light on the disk and lighting the check light on the 7133 subsystem containing the disk. Use this option to ensure the correct disk will be certified.

Note that this option toggles identify mode, but leaving the certify service aid will always reset any outstanding identify operations.

b. Certify

Select this option to certify the selected disk. Certify time depends upon the size of the disk:

- 4.5 GB Approximately ten minutes
- 2.2 GB Approximately five minutes
- 1.1 GB Approximately three minutes

Pressing the **Enter** key when certify is selected produces the following screen:

```

CERTIFY DISK 802406

Certifying:-
  pdisk1    5AEA1061  00-03-P  4GB SSA C Physical Disk Drive

                          Certify is [ 0 ] % Complete

                          [                               ]

Please stand by.

F3=Cancel      F10=Exit

```

On completion, the Certify Service Aid will indicate success or failure. The percentage complete display indicates the progress of the process as does a status bar that grows across the screen.

3. Repeatedly press the **Enter** key until the SSA Service AIDS menu is displayed.

5.1.3 Locating Devices without the Service Aids

Situations may occur where the Service Aids are not available to locate a failing device. If this should occur and Port and SSA address values have been supplied via an SRN, the following procedure can be used to locate the device.

Port and SSA addresses are represented by three digits in the SRN itself. Examples of SRNs with this information are the following:

- 43paa
- 44paa
- 45paa

Where paa is the port (p) and address (aa) respectively. The port value can be from 0 to 3 and represents the port on the SSA adapter to which the device is connected:

0	Port A1
1	Port A2
2	Port B1
3	Port B2

The SSA address is a decimal number representing the position of the device on the SSA loop to which it is connected. The first device on the loop from the port is device 0, the second is device 1, the third is device 2, and so on. Only real devices are included in the count, dummy disk modules are not counted.

Thus, in order to locate a device, locate the correct port and then count devices along the cables and through groups in the 7133 until the address value is reached. Referring to the diagrams used to design the subsystem will help with this process. The examples provided in Appendix B, "7133 Configuration Planning Charts" on page 161 can be used for this purpose.

5.2 Problem Determination and Resolution Examples

This section will provide examples of various potential problems and demonstrates how they may be detected and resolved.

5.2.1 Service Aid Errors

There are a number of errors that may occur when using the Service Aids themselves, and this section lists those errors and what to do if they occur.

SSA01 This error indicates that the host system has run out of available memory for the Service Aids to use. One of the following suggestions may resolve the problem:

- Terminate any failed application programs on the host system and then try the service aids again.
- Select the **Problem Determination** option from the diagnostics main menu and test the host system. Follow the instructions to resolve any system problems, and then try the Service Aids again.

- Shutdown and reboot the system (shutdown -Fr), and then try the Service Aids again.
- Run diagnostics from diskette or CD-ROM and use the problem determination procedures to attempt to isolate and resolve the problem, and then try the Service Aids again.

If none of these steps are successful, contact an IBM Service Representative for further assistance.

SSA02 This error indicates that an unknown problem has occurred. One of the following suggestions may resolve the problem:

- Select the **Problem Determination** option from the diagnostics main menu, and test the host system. Follow the instructions to resolve any system problems, and then try the Service Aids again.
- Run diagnostics from diskette or CD-ROM, and use the problem determination procedures to attempt to isolate and resolve the problem, and then try the Service Aids again.

If none of these steps are successful, contact an IBM service representative for further assistance.

SSA03 This error indicates that the service aids were unable to open an hdisk device. This can occur if a physical drive fails or has been removed from the subsystem. Use the following steps to resolve the problem:

1. Use the Configuration Verification service aid (see 5.1.2.3, "Configuration Verification" on page 85) to determine the location code of the adapter to which the hdisk is attached. If the location code of the hdisk is 00-03-L, for example, then the location code of the SSA adapter is 00-03.
2. Run the Link Verification service aid (see 5.1.2.2, "Link Verification" on page 82), and select this adapter.
3. If a link failure is indicated (with ????? shown somewhere in the loop), use the procedures described in 5.2.2, "Using the Service Aids to Resolve Problems" on page 94 to resolve the problem.
4. If no link failure is indicated, select the **Diagnostic Routines** option from the diagnostics menu, and then select the **System Verification** option. A screen listing all of the resources known to the host system will be displayed. Use the cursor arrow keys to move down the list (there may be several pages), and highlight the first SSA pdisk resource shown. Press the **Enter** key to test the pdisk. If no trouble was found, select the **NO** option in reply to the certify disk question, and then press the **Enter** key. Press the **Enter** key again to return to the resource list.

Repeat this process for each SSA pdisk listed. If any problems are found, follow the instructions given to resolve them.

If this procedure is unsuccessful, contact an IBM service representative for further assistance.

This completes the list of possible error messages from the Service Aids themselves.

5.2.2 Using the Service Aids to Resolve Problems

This section will look at how the Service Aids can be used to resolve any problems that may occur.

5.2.2.1 Problem Determination Procedures

The following process should be followed to detect and analyze problems:

1. POSTs.

Power-On Self-Test operations run in each disk drive and in every SSA adapter at power-up time. These tests are responsible for checking that the hardware and microcode are within the device are operating correctly. Any errors detected during a POST will be reported via either a check light on the device (disks) or error messages in the system log (SSA adapter).

2. Monitor the host system error log.

Any errors detected by the 7133 subsystem, four-port adapter or software will be appended to the system error log on the host system. These error messages will contain the following information:

SRN The Service Reference Number identifying the error.

FRU list The list of Field Replaceable Units that may be causing the problem.

Percentage The percentage chance that a particular FRU is causing the problem.

Problem A description of the problem and the action that must be taken.

3. Attempt to solve the problem by taking the recommended action.

This may involve referring to the *7133 Service Guide* or other manuals supplied with the 7133 and host system.

4. Use the information in this section to assist the process.

Use the information in the rest of this section as a supplement or for basic problem determination assistance.

Important

Never exchange parts unless instructed to do so by the *7133 Service Guide* or by a qualified IBM Service Representative.

5. Call an IBM Service Representative

If none of the above procedures are effective, call an IBM service representative to assist in solving the problem.

5.2.2.2 Problem Solving Procedures

There are various kinds problems that may occur and these may be categorized as follows:

- Software problems

Problems with any of the software elements that collectively support the hardware including diagnostics, device drivers and microcode.

- Link problems

Problems related to physical connectivity.

- Adapter problems

Problems related to the SSA four port adapter hardware.

- Disk problems

Physical disk related problems.

- Hardware problems

Problems with associated hardware such as power and cooling modules.

Each of these areas will now be discussed in more detail.

Software Problems: Software problems can occur for any of the following reasons:

- Incorrect levels of software
- Invalid configurations
- Software bugs

If problem determination indicates that the software is the probable cause, then check the following:

1. Incorrect levels of software

Confirm that the latest level of the SSA support software is loaded, Use the information in Chapter 3, "Installation and Configuration" on page 33 to check this.

2. Invalid configurations

It is possible that the software has been incorrectly installed and configured. Check the error logs to confirm that installation was successful. If not, reinstallation may be necessary.

3. Software bugs

Certain symptoms may require PTFs to clear up. If unexplainable in any other fashion, report the problem to an IBM Service Representative.

If problems still persist, then the software is unlikely to be causing them. Return to problem determination.

Link Problems: Link problems can occur for any of the following reasons:

- Invalid configurations
- Cables are accidentally disconnected
- Cables become damaged

If problem determination indicates links as the probable cause, then check for the following:

1. Invalid configurations

Use the information in Chapter 3, "Installation and Configuration" on page 33 to check that a valid SSA configuration has been created. In particular, check for the following:

- Open loops
- More than 48 physical drives on a loop
- More than eight adapters on a loop
- More than one adapter port pair per adapter on a loop

- Cable lengths exceeding the maximum allowable distance

If the configuration is valid, check the cables themselves as described next.

2. Disconnected or damaged cables

Use the information in 5.1.2.2, “Link Verification” on page 82 to check whether there is a failure in the loop. If a failure has been detected, isolate the physical disk in the loop where the failure has occurred, and check that the cables on either side of this disk are connected securely.

If they are and the problem still persists, try replacing the cables.

If this does not solve the problem, check the disk drives themselves as described in “Disk Problems.”

If the problem still exists then a link fault is not causing it. Return to problem determination.

Adapter Problems: Adapter problems can occur for any of the following reasons:

- Poorly seated adapter
- Incorrect microcode level
- Incorrect adapter level
- Damaged adapter

If problem determination indicates that the adapter is the probable cause, then check the following:

1. Adapter seating

Ensure that the adapter card is physically seated correctly in its slot.

2. Microcode/adapter levels

Check that the levels of microcode and the card level are compatible with each other and with the operating system level. Refer to the information in Chapter 3, “Installation and Configuration” on page 33 to confirm this.

3. Damaged adapter

It is possible that the adapter itself may have failed. Try replacing the adapter.

If problems still persist, then the adapter is unlikely to be causing them. Return to problem determination.

Disk Problems: Disk problems can occur for any of the following reasons:

- Incorrect disk drive type
- Poorly seated drive
- Drive communications errors
- Drive media errors

If problem determination indicates that a drive is the probable cause, then check for the following:

1. Incorrect drive type

Ensure that only supported SSA disk drives have been installed in the 7133 subsystems. The only drives that are currently supported are the IBM SSA disk drives:

- 4.5 GB - 88G6401
- 2.2 GB - 88G6400
- 1.1 GB - 07H8230

If there are any drives not bearing these part numbers in a loop, remove them, and replace them with the correct drives.

2. Poorly seated drive/drive removed

It is possible that a drive may have been poorly inserted into a subsystem slot or that a drive has been removed. Use the information in 5.1.2.2, "Link Verification" on page 82 to determine which drive may have the problem, and then check to see if the drive is missing or not inserted properly.

3. Communications errors

The communications electronics on the disk may be damaged, preventing the disk from properly attaching to the loop. Try replacing the disk drive. If the new drive functions correctly, then the old drive will need servicing or permanent replacement.

4. Media errors

There may be media errors on the disk. Use the information detailed in 5.1.2.5, "Certify Disk" on page 90 to certify the disk. If this procedure shows errors, the disk may need replacing. Back up any information on the disk, and use the information in 5.1.2.4, "Format Disk" on page 87 to reformat the disk. If excessive media errors still occur, the disk will need replacing.

If problems still persist, then disk drives are unlikely to be causing them. Return to problem determination.

Hardware Problems: Hardware problems may occur for any of the following reasons:

- Damaged power supplies
- Damaged remote power control feature
- Damaged cooling system
- Damaged signal cards
- Damaged cables

If problem determination indicates that any of these possibilities have occurred, then the parts in question will probably need replacing. Use the MAPs (Maintenance Analysis Procedures) in the *7133 SSA Disk Subsystem Service Guide* to solve the problem.

Chapter 6. Performance/Availability Investigation

This chapter investigates the actual effect of the various performance optimization methods discussed in Chapter 4, "Hardware Scenarios" on page 49. In addition, the effect of tuning those SSA parameters that can be changed is examined.

The host system used for the tests was an RS/6000 model 580 with 96 MB of memory running AIX V3.2.5 and containing a single SSA four-port adapter.

Disclaimer

Please note that these tests are not intended to provide definitive performance information. The purpose of this chapter is to investigate the relative effect of implementing various strategies to improve performance.

These figures should not be considered representative of the maximum or minimum performance capabilities of the 7133 SSA subsystem. The environment in which they were measured was specific to these tests and therefore inappropriate for general usage.

6.1 Bandwidth

This section will look at the bandwidth utilization of an SSA loop. The maximum sustainable throughput to individual disks, and for the loop itself, will be tested. The I/O generated will be varied as follows:

- Sequential/random

The I/O generated to each disk will be varied between sequential and random reads and writes. Sequential I/O reads and writes consecutive blocks, while random I/O reads and writes to blocks scattered randomly throughout a file.

- Block size

The size of the individual data blocks read and written will be varied.

- Read/write ratio

The ratio of the number of read operations to the number of write operations being executed will be varied.

- File size

The size of the files created within which reads and writes are performed will be varied.

- Number of processes

The number of processes actually performing the I/O will be varied.

Tables of the average throughput while these parameters are varied will be provided.

6.1.1 Individual Disk

This section will look at the maximum throughput to an individual disk. The disk being used is the second one from port A1 in a loop of 10 disks attached between port A1 and port A2 on the same adapter. There is no other activity to the other disks on the loop.

A single process performs I/O to the disk, while the read/write ratio, block size, file size, and I/O type are varied. The performance indicators measured are the average utilization of the disk over the time of each test, the average throughput to the disk and the average number of transactions per second.

Table 4. Maximum Throughput to a Single Disk. The maximum sustained throughput to an individual disk as several I/O parameters are varied.

I/O Parameters				Average Percentage Utilization	Average Throughput (KB)	Average Transactions Per Second
File Size	R/W Ratio	Block Size	I/O Type			
10 K	75	4 K	R	0.03	0.20	0.05
50 K	75	4 K	R	79.85	1129.08	140.60
100 K	75	4 K	R	82.25	1126.89	140.38
1 M	75	4 K	R	93.47	1149.71	143.11
10 M	75	4 K	R	98.70	1026.41	123.32
500 K	25	4 K	R	97.18	1297.59	158.13
500 K	50	4 K	R	95.55	1257.88	155.94
500 K	75	4 K	R	91.10	1179.30	146.09
500 K	75	4 K	S	98.13	2945.73	116.47
500 K	75	8 K	R	90.74	1496.61	123.47
500 K	75	16 K	R	88.55	1886.22	114.02
500 K	75	32 K	R	88.82	2287.03	96.72
500 K	75	64 K	R	83.49	3327.42	64.13

The following facts are evident from Table 4:

1. File size

In this series of tests, the file size was increased while the block size was kept at 4 KB, 75 percent of operations were reads and all I/O was random.

At a file size of 10 KB, most of the data can be kept in memory so the utilization, throughput and transactions per second are all extremely low.

As the file size increases, the utilization of the disk increases as would be expected and longer seeks need to be performed. Average throughput and transactions per second stay constant at around 1.1 MB per second and 140 TPS up to a file size of 10 MB. At this point, the throughput and transactions per second fall off slightly, indicating that the seek time has increased.

2. Read/write ratio

In this series of tests, the read/write ratio was varied while the file size was kept at 500 KB, the block size kept at 4 KB and the I/O maintained random.

The utilization of the disk falls off as the ratio of reads increases, as do the throughput and transactions per second. It would be expected that the utilization of the disk would increase with a higher percentage of writes, and this is the case. It is most likely that the throughput and transactions per second fall as more reading is performed due to the effect of caching. The more reads that are performed, the greater the likelihood that a read page will already be in system memory, particularly with the relatively small file size used.

3. Sequential/random I/O

The type of I/O was switched from random to sequential for a file size of 500 KB, a block size of 4 KB and with a read/write ratio of 75 percent reads.

The utilization increased as would be expected, since sequential I/O is more efficient. The throughput more than doubled, again as expected, as disk seeking is minimized for sequential operations. The transactions per second fell again, most likely due to the effect of file system read ahead. When the system detects sequential reading, it begins to read ahead in increasing quantities, thus reducing the number of I/O operations required while increasing the throughput.

4. Block size

The block size was increased while the file size was kept at 500 KB, the read/write ratio at 75 percent reads and the type of I/O random.

The utilization decreases steadily as is expected. More time is spent performing the reading and writing of data than the more time consuming seeks, thus utilization falls. Throughput, however, increases markedly, again as expected. More time is spent actually streaming data to and from the disk, so the throughput increases. The transactions per second decrease. As each I/O operation now reads or writes increasing amounts of data, the time taken increases, and consequently, the number of transactions per second decreases.

6.1.2 SSA Loop

This section will look at the maximum throughput for the loop itself. The number of disks used will be increased from 1 to 10 disks. Again the loop contains 10 disks and is connected between ports A1 and A2 on the adapter. While disks are not being used, they remain attached to the loop.

One process is active per disk that is being used and each process performs identical I/O to a 10 MB file. The data is read/written sequentially in blocks of 96 KB with 10 percent of the total I/O being read operations. The indicators measured are the average utilization of each disk over the time of each test, the total throughput to the loop and the average number of transactions per second for the loop.

Table 5. Maximum Throughput to an SSA Loop. The maximum sustained throughput to an SSA loop as several I/O parameters and the number of disks are varied.

Number of Disks	Processes	Average Percentage Utilization	Total Throughput (KB)	Total Transactions
1	1	100.00	6627.86	104.72
2	2	98.06	11982.07	195.20
3	3	98.12	16083.47	267.79
4	4	97.96	17016.44	279.72
5	5	78.42	19777.41	312.80
6	6	76.83	19448.21	328.64
7	7	62.88	19486.63	318.01
8	8	57.17	19045.30	382.56
9	9	48.18	18560.85	380.03
10	10	42.36	18745.41	338.85

The following facts are evident from Table 5:

As the number of disks increases, the average utilization of each disk remains reasonably constant at the maximum, up to four disks. As more disks are added, the utilization now falls off rapidly, indicating that the disks are not now receiving enough transactions to operate at full utilization.

From the figures for one disk, it can be seen that the maximum throughput to a single disk is around 6.6 MB per second. As the number of disks is increased, the total throughput on the loop increases to around 19.7 MB per second and then levels off. In this environment, the throughput is being constrained by the CPU performance rather than loop performance; the actual maximums could be higher, but it can be seen that the knee of the curve is at around four disks in this sequential environment.

Again, from the figures for a single disk, it can be seen that this I/O configuration generates around 100 transactions per second per disk. As the number of disks increases, the transactions do not rise linearly; this is not due to any constraints as the transaction rate for the adapter is nowhere near, but due to the fall off in the throughput caused by the CPU constraint.

Thus, it can be seen that the total number of disks that should be supported on a loop depends upon the I/O profile and in particular, on whether the I/O is sequential or random. For large block size sequential I/O, the knee of the performance curve is at around four disks.

6.2 Data Path

This section will look at the effect of the location of a disk in the loop on throughput. The average throughput for each disk on a loop that is being fully utilized will be monitored to ascertain which disk(s) show the maximum throughput.

One process is active per disk that is being used and each process performs identical I/O to a 10 MB file. The data is read/written sequentially in blocks of 96 KB, with 10 percent of the total I/O being read operations. The indicators measured are the average utilization of each disk over the time of each test, the total throughput to the loop and the average number of transactions per second for the loop.

The tests are performed initially with 10 disks active, then with nine, and finally with eight active disks.

<i>Table 6. Effect of Disk Position on Throughput (10 disks). The variance in average throughput with the position of the disk in a 10-disk SSA loop.</i>			
Disk Position	Average Percentage Utilization	Average Throughput (KB)	Average Transactions
0	58.12	1898.47	33.32
1	45.12	1865.34	34.06
2	37.87	1886.07	34.46
3	56.49	1941.17	34.99
4	36.21	1896.00	35.44
5	51.52	1848.03	32.41
6	36.38	1947.60	37.96
7	33.48	1801.43	31.21
8	34.47	1850.60	32.76
9	33.51	1810.70	32.24

From the results in Table 6, it can be seen that the average throughput to each disk is fairly uniform, as are the number of transactions per second. The first disk in the loop does seem to demonstrate a slightly higher utilization than the others, though there are peaks in the throughput and transaction rate for other disks too; for example, disk 6 has a slightly higher throughput and transaction rate than the other disks.

Nothing conclusive can be established which implies that the position of the disk in a loop is not significant with this sequential I/O profile.

Table 7. Effect of Disk Position on Throughput (nine disks). The variance in average throughput with the position of the disk in a nine-disk SSA loop.

Disk Position	Average Percentage Utilization	Average Throughput (KB)	Average Transactions
0	67.68	2154.90	43.41
2	40.92	1974.45	38.93
3	65.53	2191.14	45.86
4	40.17	2031.88	41.24
5	63.83	2199.22	46.71
6	38.47	2016.16	39.91
7	38.59	1972.68	38.64
8	38.96	2008.10	41.17
9	39.51	1972.68	38.64

The results in Table 7 show the same measurements as for Table 6 on page 103, except with nine disks used instead of 10. Once again, disk 0 has the highest utilization, and again there are peaks in the throughput and transactions for other disks. Those disks exhibiting the highest transaction rate and throughput are different in this test than for the previous, which reinforces the fact that there is no pattern emerging for maximizing throughput based on disk position.

Table 8. Effect of Disk Position on Throughput (eight disks). The variance in average throughput with the position of the disk in a eight-disk SSA loop.

Disk Position	Average Percentage Utilization	Average Throughput (KB)	Average Transactions
0	82.60	2636.50	52.34
3	69.08	2325.40	49.47
4	44.06	2223.23	46.09
5	77.19	2654.83	53.85
6	46.13	2294.85	46.31
7	47.32	2345.21	46.07
8	46.70	2360.31	42.56
9	44.27	2345.21	46.07

The results in Table 8 show a similar pattern to the previous tests, except this time with eight disks utilized in the loop. Disk 0 has the highest utilization, and again there are peaks in the throughput and transactions for other disks. Those disks exhibiting the highest transaction rate and throughput are again different in this test than for the previous, which reinforces the fact that there is no real pattern emerging for maximizing throughput based on disk position.

6.3 Multiple Adapters

This section looks at the effect on performance of having multiple adapters on a loop. The throughput to a fully utilized loop between adapters in two separate host systems is monitored.

One process is active per disk that is being used, and each process performs identical I/O to a 10 MB file. The data is read/written sequentially in blocks of 96 KB, with 10 percent of the total I/O being read operations. The indicators measured are the average utilization of each disk over the time of each test, the total throughput to the loop and the average number of transactions per second for the loop.

For this, two adapters were used, each in a separate host system. Two loops were connected between the two adapters, a loop with two disks on between port A2 on adapter 1 and port A2 on adapter 2 and a loop with eight disks on between port A1 on adapter 1 and port A1 on adapter 2. The I/O generated is initially to the disks on the eight-disk loop, but as the number of disks increases beyond eight, includes the disks on the other loop.

I/O is only generated from one of the systems.

Table 9. Effect of Multiple Adapters on Throughput. The variance in average throughput with the number of adapters on the SSA loop.

Number of Disks	Average Percentage Utilization	Total Throughput (KB)	Total Transactions
5	97.16	18130.63	343.04
6	95.57	19436.08	409.70
7	77.28	20179.39	395.02
8	59.26	19510.42	350.89
9	49.19	19046.23	347.50
10	39.75	17914.71	337.79

From the results in Table 9, it can be seen that as would be expected, the average utilization of the disks falls as the number of disks used increases. The throughput rises to a maximum of around 20 MB per second for seven disks and then gradually falls off to 17 MB per second for 10 disks. The transaction rate rises to a maximum of 409.7 for six disks and then gradually falls off as the number of disks increases.

Comparing this output to the results in Table 5 on page 102 shows very similar results. It appears from this test that a second adapter does not affect the overall performance. In this test, the second adapter was in a second host system that was generating no I/O, so this test really illustrates that the presence of a second adapter does not have any significant impact on performance.

6.4 Data Location

This section looks at the effect that the location of data on the physical disk itself has on performance. The maximum throughput to file systems on the outer edge and center of a disk will be compared.

Initially in this test, two processes are active, one writing to the center of a disk and one to the edge of the same disk. In the second part of the test, four processes are active, two writing to the center of a disk and two to the edge of the same disk. The I/O is random to 5 MB files, with 50 percent reads and a block size of 4 KB. The performance indicators measured are the number of reads and writes made by each process over the period of the test.

Table 10. Effect of Data Location on Throughput. The variance in average throughput with the location of the data on a physical disk.

Number of processes	Process	Data Location	Number of reads	Number of writes
2	0	Center	15318	15587
2	1	Edge	14981	15368
4	0	Center	8974	9037
4	1	Center	7984	8149
4	2	Edge	8211	8217
4	3	Edge	8830	8911

From the results in Table 10, it can be seen that in the first part of the test, with a single process writing to the center of the disk and a second writing to the edge of the same disk, that there is no significant difference in the throughput of either process.

When the number of processes writing is doubled, the amount of I/O per process decreases, but again there is no significant difference in throughput between the center and the edge of the disk.

From this, it can be concluded that there is no physical reason for performance improvement based on data location, but the statistical likelihood that the disk head will be over the center of the disk more often for a disk whose surface is uniformly utilized still holds true. Thus the most highly accessed data should be placed in the center of the disk for optimum performance in a uniformly utilized disk.

6.5 Multiple Disk Types

This section will investigate the effect of concurrency of access on performance. The difference in throughput for several processes reading and writing to files on a single disk and those same processes reading and writing to multiple disks will be compared.

In the first part of this test, 10 processes access small files on a single disk and then 10 processes access large files on the disk. In the second part of the test five disks are used with two processes accessing each disk. Again, first small files are used and then large files. In all cases, I/O is random using a block size

of 4 KB and 10 percent reads; the large files are 10 MB and the small files 100 KB. The performance indicators measured are the average disk utilization, the total throughput and the total transactions per second for each disk used.

Table 11. Effect of Concurrency on Throughput. The variance in average throughput with concurrency of access to the data.

Number of Disks	Number of Processes per Disk	File Size	Average Percentage Utilization	Total Throughput (KB)	Total Transactions
1	10	Small	98.18	1253.31	141.52
1	10	Large	98.32	1991.26	83.77
5	2	Small	97.50	6256.19	750.62
5	2	Large	93.59	5594.88	623.73

From the results in Table 11, it can be seen that by spreading the files across more disks, the concurrency thus made possible increases throughput and the number of transactions per second. The relative increase is slightly higher for larger files than for smaller, though in both cases, increasing the number of disks to five has increased performance by approximately five.

6.6 Tuning SSA Parameters

The main parameter that can be changed is the `queue_depth` for the `hdisks`. In order to change the queue depth for an `hdisk`, use the following procedure:

```
# umount /filesystem
# chpv -vr hdiskn
# chdev -l hdiskn -a "queue_depth=x"
hdiskn changed
# chpv -va hdiskn
# mount /filesystem
```

Any file systems on the `hdisk` must be unmounted and the `hdisk` made unavailable with the `chpv -vr hdiskn` command, where `n` is the number of the `hdisk` to be changed. The `chdev` command can then be used to alter the value of the queue depth as shown. The `hdisk` must then be made available using the `chpv -va` command and the file systems mounted again.

The current value of the queue depth for an `hdisk` can be displayed using the following command:

```
# lsattr -El hdiskn -a queue_depth
queue_depth 3 Queue depth True
```

Where `n` is the number of the `hdisk` to be displayed.

Note

It is not recommended to change the queue depth parameter as the system default is probably optimum for most environments.

Appendix A. SSA Product Reference Information

This chapter provides detailed product specific information on the IBM 7133 Serial Storage Architecture Disk Subsystem. The various options and features available for both models of the 7133 are listed for reference purposes.

A.1 IBM 7133 SSA Disk Subsystem Model 010

The 7133 provides power and packaging for SSA disk drive modules. The model 010 is designed to be mounted in a standard 19-inch rack, and is four EIA units high. The model 010 provides slots for up to 16 SSA disk drive modules mounted in convenient carriers. These disk drive modules are auto-docking and may be replaced by the customer. The drives connect to four internal SSA loops which can be configured as independent loops or linked via external cables.

The 7133 can be populated with three different SSA disk drive modules. The capacities supported are 1.1GB, 2.2GB and 4.5GB. The machine includes four 2.2GB disk drives as standard. These can be replaced by eight 1.1GB disk drive modules or four 4.5GB drives at the time of the original order.

Features are provided to add any combinations of supported drives. The 7133 includes two auto-docking power/cooling units. These will provide redundant power and cooling to disk drive module positions 1 to 4. They will also provide non-redundant power/cooling to positions 5 to 8. An additional unit is available as a selective feature, and this will provide redundant power and cooling to all 16 disk drive module positions. These units may also be replaced by the customer.

The 7133 has selectable cable features, up to 25 meters long, for attaching to an IBM High Performance 4-Port SSA Adapter in the attaching system, for linking the internal busses, and for connecting to other 7133 subsystems.

A.1.1 Physical Specifications

The 7133 model 010 has the following physical specifications:

- Width: 444mm (17.5 inches)
- Depth: 665mm (26.2 inches)
- Height: 171mm (6.7 inches)
- Weight: 32Kg (70 lbs) Entry configuration
- Weight: 50Kg (110 lbs) Maximum configuration

A.1.2 Operating Environment

The 7133 model 010 requires the following operating environment:

- Temperature: 10 to 40 degrees C
- Relative Humidity: 8 to 80%, non-condensing.
- Wet Bulb: 27 degrees C Calorific value: 1810 BTU/hr (maximum config)
- Electrical power : 0.645KVA (maximum config, start up)
- Electrical power : 0.480KVA (maximum config, operating)

- Electrical power : 0.420KVA (maximum config, idling)
- Capacity of Exhaust : 25 litres/second
- Noise Level : 5.5Bels (maximum config, idling)
- Noise Level : 5.6Bels (maximum config, operating)

The EMC Conformance Classifications:

- USA - FCC Class A
- Germany - VDE Class A
- Japan - VCC-1

The product safety/country testing/certification:

- USA - UL
- Canada - CSA
- Germany - GS Mark (Safety, TUV)
- International - IEC

A.1.3 Limitations

Not applicable.

A.1.4 Hardware Requirements

The 7133 Model 010 attaches to the following RISC System/6000 and POWERparallel systems:

- 7012 Models 34H, 340, 350, 360, 370, 380, 390, and 39H.
- 7013 Models 52H, 53E, 53H, 55E, 55S, 55L, 550, 56F, 560, 57F, 570, 58F, 58H, 580, 59H, 591, and 590.
- 7015 Models R10, R20, R21, R24, 95E, 950, 97E, 970, 97F, 97B, 980, 98E, 98F, 98B, 99E, 99F, 99J, 99K and 990.
- 9076 Models 201, 202, 203, 204, 301, 302, 303, 304, 401, 402, 403, 404, 3A2, 3A3, 3A4, 3B2, 3B3, and 3B4.

The attaching system must provide 4 EIA units of rack space for the 7133 in the system rack, where applicable, or in a 7202 model 900 expansion Rack. It must also provide an attachment to an SSA loop via feature code 6214, the SSA adapter.

A.1.5 Software Requirements

The IBM 7133 SSA Disk Subsystem is supported by AIX V3.2.5 with additional PTFs, and the AIX 3.2.5 device driver shipped with the SSA Four Port Adapter (feature code 6214 on the attaching system). For ease of installation these PTFs are packaged with the Device Driver on the CD-ROM shipped with the adapter. Customers without access to CD-ROM drives on their machines or network can obtain the driver and the required PTFs via FIXDIST. The device driver is available as APAR IX52018. The required PTFs, on FIXDIST, are identified as PMP3251. For alternative delivery, contact your Software Service representatives for the appropriate PTFs. The additional Version 3.2.5 PTFs (without the AIX V3.2.5 Device Driver for the adapter) are included on all AIX V3.2.5 orders shipped after May 19, 1995 labelled AIX 3.2.5 Enhancement 5 (3250-05-00).

A.1.6 Electronic Customer Support

Not available.

A.1.7 Publications

The publications listed in Table 12 are shipped with the product. Additional copies are available immediately.

<i>Table 12. Publications shipped with the 7133</i>	
Title	Order Number
Operator Guide	GA33-3259
Installation Guide (Model 010)	GA33-3260
Service Guide	SY33-0185
Note: Only the Ship Pack publications appear in the Sales Manual. A complete list of publications is provided in the PUBS section of HONE.	

A.1.8 Features

The following features are available for the 7133 model 010:

A.1.8.1 No Charge Specify Codes

The features listed in Table 13 summarize the no charge specify codes for the 7133 model 010.

<i>Table 13. No charge feature code summary</i>	
Description	Feature
Havant Manufacturing Direct Ship to Customers	0550
RSD Austin Manufacturing Integration	0987
Attachment System - RISC System Model 7012	7012
Attachment System - RISC System Model 7013	7013
Attachment System - RISC System Model 7015	7015
Attachment System - RS/6000 SP Model 9076	9076
Standard 2.2GB SSA disk drive modules, one of four	9204
100/200V 50/60Hz and 300V DC input power	9850
Language Groups - US English	9300
Language Groups - French	9703
Language Groups - German	9704
Language Groups - Spanish	9708
Language Groups - Italian	9711
Language Groups - Japanese	9714

The following list provides details of the no charge specify codes:

SC 9850 100/200V 50/60HZ and 300VDC input power

This feature supplies two auto-sensing power and cooling modules that will operate with input ac power between 100 and 200V, or 300V DC. This feature is mandatory for all orders.

SC 9204 Standard 2.2GB SSA disk drive modules, one of four

This feature provides 2.2GB SSA disk drive modules mounted in positions 1 to 4. They are included as standard, and may be replaced by 1.1GB or 4.5GB SSA disk drive modules at time of original order

only. The specification of the disk drive modules is as described under feature code 3201. This code will be specified four times if none of the disk drive select features are ordered.

SC 0550 Havant manufacturing direct ship to customers

This specify code is for administrative purposes only. This code must be added to all orders for the IBM 7133 Model 010 SSA Disk Subsystem that are to be shipped directly from Havant to customer locations, and are not tied to a new system order. All orders must specify either this code, feature code 0003 or feature code 0987.

SC 0987 RSD Austin manufacturing integration

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 Model 010 SSA Disk Subsystem with new system orders. This will ensure that the product is integrated and shipped with the system. All orders must specify either this code, feature code 0003 or feature code 0550.

SC 9076 System attachment - SP2 POWERParallel System

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 SSA Disk Subsystem with new SP2 system orders. This will ensure that the product will be recorded against the intended RS/6000 SP or RS/6000 system on which it will be installed. All orders must specify either this code, feature code 7015 or feature code 7012 or feature code 7013.

SC 7015 System attachment - RISC System/6000 Model 7015

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 SSA Disk Subsystem with new RISC system/6000 7015 systems orders. This will ensure that the product will be recorded against the intended RS/6000 SP or RS/6000 system on which it will be installed. All orders must specify either this code, feature code 9076 or feature code 7012 or feature code 7013.

SC 7012 System attachment - RISC System/6000 Model 7012

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 SSA Disk Subsystem with new RISC System/6000 7012 systems orders. This will ensure that the product will be recorded against the intended RS/6000 SP or RS/6000 system on which it will be installed. All orders must specify either this code, feature code 9076 or feature code 7015 or feature code 7013.

SC 7013 System attachment - RISC System/6000 Model 7013

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 SSA Disk Subsystem with new RISC System/6000 7013 systems orders. This will ensure that the product will be recorded against the intended RS/600 SP or RS/6000 system on which it will be installed. All orders must specify either this code, feature code 9076 or feature code 7015 or feature code 7012.

A.1.8.2 Chargeable Special Feature Codes

The features listed in Table 14 summarize the chargeable specify feature codes for the 7133 model 010.

Description	Feature
Remote power-on control	3001
Power/cooling feature 3	3003
Disk Drive Module Lock	3004
Rack Blanking Plates	3005
One 1.1GB SSA disk drive module	3101
1.1GB SSA disk drive module, one of eight (select)	3108
.1GB disk drive module, one of sixteen (select)	3116
One 2.2GB SSA disk drive module	3201
2.2GB SSA disk drive module, one of four	3208
2.2GB SSA disk drive module, one of twelve	3216
One 4.5GB SSA disk drive module	3401
4.5GB SSA disk drive module, one of four (select)	3404
4.5GB SSA disk drive module, one of eight (select)	3408
4.5GB SSA disk drive module, one of sixteen (select)	3416
0.18M SSA copper cable	5002
0.60M SSA copper cable	5006
1.00M SSA copper cable	5010
2.50M SSA copper cable	5025
5.00M SSA copper cable	5050
10.0M SSA copper cable	5100
25.0M SSA copper cable	5250
2.50M SSA Subsystem to Additional System Cable	5604
5.00M SSA Subsystem to Additional System Cable	5605
10.0M SSA Subsystem to Additional System Cable	5606
25.0M SSA Subsystem to Additional System Cable	5607
2.50M SSA Subsystem to Subsystem Cable	5704
5.00M SSA Subsystem to Subsystem Cable	5705
10.0M SSA Subsystem to Subsystem Cable	5706
25.0M SSA Subsystem to Subsystem Cable	5707
0.18M SSA Loop to Loop Cable	5801
1.00M SSA Loop to Loop Cable	5803
2.50M SSA Subsystem to Adapter Cable	5904
5.00M SSA Subsystem to Adapter Cable	5905
10.0M SSA Subsystem to Adapter Cable	5906
25.0M SSA Subsystem to Adapter Cable	5907
The following features must be ordered against the 7132-REP machine	
1.1GB SSA disk drive module replacement	3199
2.2GB SSA disk drive module replacement	3299
4.5GB SSA disk drive module replacement	3499
Power/cooling replacement unit	3099
Disk Drive Module Lock	3004

The following list provides details of the chargeable specify codes:

FC 3001 REMOTE POWER-ON CONTROL

This feature allows the attaching system(s) to control the power on/off status of the 7133 SSA Disk Subsystem.

- Technical Description

The feature will control the power on/off status of the power supplies within the 7133 SSA Disk Subsystem under control of up to two systems. When two systems are connected, the product

will power on if requested by either host, and will power off when requested by both.

- Maximum features per model 010

One.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

None.

FC 3003 POWER/COOLING FEATURE 3

This feature provides redundant power and cooling to disk drive positions 5 to 16.

- Technical Description

The 7133 SSA Subsystem is supplied with two auto-docking combined power/cooling units. These supply redundant power/cooling to disk drive positions 1 to 4, and non-redundant power/cooling to positions 5 to 8. This feature is mandatory for disk drive positions 9 to 16. This feature provides a third auto-docking unit which supplies redundant power/cooling to the complete subsystem.

- Maximum features per model 010

One.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable
Yes.
- Cable order
None.

FC 3004 DISK DRIVE MODULE LOCK

This feature provides a lock which can be used to control removal of the disk drive modules. The first order of this feature will provide three keys and a lock. Subsequent orders in the same total order will provide only a lock.

- Technical Description

Disk drive modules are fitted with a lever which ensures correct connection of the module on insertion, and controlled removal when required. The lock supplied by this feature fits on the front of the module and inhibits the use of the lever. There is only one version of the lock. This feature prevents inadvertent removal of the module. It cannot be considered as protection against malicious removal.

- Maximum features per model 010
Sixteen.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
None.

FC 3005 RACK BLANKING PLATES

This feature provides blanking plates to cover 4 EIA units of rack space in a 19-inch rack. It should be ordered on installed systems when 7133 model 010 installations displace other storage products, resulting in unused rack space.

- Technical Description

The feature consists of two flat cosmetic plates with mounting hardware for attaching to a 19-inch rack. These are the same plates used in the original rack manufacturing.

- Maximum features per model 010

- Five.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
None.

FC 3099 POWER/COOLING REPLACEMENT UNIT

This feature supplies an on-site spare power/cooling unit which a customer can use to rapidly replace a power/cooling unit in the unlikely event of a failure.

- Technical Description
The hardware supplied by this feature is the same as provided by feature code 3003, but includes protective packaging, and installation instructions.
- Maximum features per model 010
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
Yes
- Limitations
Must be ordered against 7132-REP with one power/cooling unit per order.
- Field installable
Yes.
- Cable order
None.

FC 3101 ONE 1.1GB SSA DISK DRIVE

This feature supplies one 1.1GB SSA disk drive mounted in an auto-docking carrier. The auto-docking module includes a lever that facilitates safe insertion or removal of the disk drive module.

- Technical Description

The feature consists of a 1.1GB 3.5 inch form factor device mounted on a carrier. The carrier houses power conversion circuits in addition to the disk drive. The drive characteristics are as described below:

- Formatted Capacity - 1.1GB
- Form Factor - 3.5inch
- Average Seek Time - 6.9 msec
- Latency - 4.17 msec
- Media Data Transfer Rate - 9.59 to 12.58 MB/sec (10 bands)
- SSA Data Transfer Rate - 20 MB/sec Full Duplex

- Maximum features per model 010

Sixteen. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

Yes.

- Cable order

None.

- Maximum order quantity

Sixteen.

FC 3108 1.1GB SSA DISK DRIVE MODULE, ONE OF EIGHT (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of eight 1.1GB drives in positions 1 to 8. The feature is priced to provide credit for the removal of the 2.2GB drives. This

feature is available at time of first order only and must be ordered eight times. Note that for redundant power for the positions 5 to 8, feature code 3003 Power/cooling feature 3 must be ordered.

- Technical Description

The feature content is the same as feature code 3101.

- Maximum features per model 010

Eight. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Eight.

FC 3116 1.1GB SSA DISK DRIVE MODULES, ONE OF SIXTEEN (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive normally included with the machine and installation of sixteen 1.1GB drives in positions 1 to 16. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered sixteen times.

- Technical Description

The feature content is the same as feature code 3101.

- Maximum features per model 010

Sixteen. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

feature code 3003, Power/cooling feature 3.

- Corequisites

- None.
- Compatibility conflicts
 - None.
- Customer setup
 - No.
- Limitations
 - The maximum number of drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features: 3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.
- Field installable
 - No.
- Cable order
 - None.
- Minimum order quantity
 - Sixteen.

FC 3199 1.1GB SSA DISK DRIVE MODULE REPLACEMENT

This feature supplies an on-site spare 1.1GB disk drive module which a customer can use to rapidly replace a disk drive module in the unlikely event of a 1.1GB drive failure.

- Technical Description
 - The hardware supplied by this feature is the same as provided by feature code 3101, but includes protective packaging, and installation instructions.
- Maximum number of features
 - One per order.
- Prerequisites
 - None.
- Corequisites
 - None.
- Compatibility conflicts
 - None.
- Customer setup
 - Yes
- Limitations
 - Must be ordered against 7132-REP with one disk drive module per order.
- Field installable
 - Yes.
- Cable order
 - None.

FC 3201 ONE 2.2GB SSA DISK DRIVE MODULE

This feature supplies one 2.2GB SSA disk drive mounted in an auto-docking carrier. The auto-docking carrier includes a lever that facilitates safe insertion or removal of the disk drive module.

- Technical Description

The feature consists of a 2.2GB 3.5-inch form factor device mounted on a carrier. The carrier houses power conversion circuits in addition to the disk drive. The drive characteristics are as described below:

- Formatted Capacity - 2.2GB
- Form Factor - 3.5 inch
- Average Seek Time - 7.5 msec
- Latency - 4.17 msec
- Media Data Transfer Rate - 9.59 to 12.58 MB/sec (10 bands)
- SSA Data Transfer Rate - 20 MB/sec Full Duplex

- Maximum features per model 010

Sixteen. Note limitations of total drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

Yes.

- Cable order

None.

FC 3208 2.2GB SSA DISK DRIVE MODULES, ONE OF FOUR

Ordering this feature results in the installation of four 2.2GB drives in positions 5 to 8. This feature is available at time of first order only and must be ordered four times. Note that for redundant power for the positions 5 to 8, feature code 3003 Power/cooling feature 3 must be ordered.

- Technical Description

The feature content is the same as feature code 3201.

- Maximum features per model 010

Four. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Four.

FC 3216 2.2GB SSA DISK DRIVE MODULES, ONE OF TWELVE

Ordering this feature results in the installation of twelve 2.2GB drives in positions 5 to 16. This feature is available at time of first order only and must be ordered twelve times.

- Technical Description

The feature content is the same as feature code 3201.

- Maximum features per model 010

Twelve. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

feature code 3003, Power/cooling feature 3.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Twelve.

FC 3299 2.2GB SSA DISK DRIVE MODULE REPLACEMENT

This feature supplies an on-site spare 2.2GB disk drive module which a customer can use to rapidly replace a disk drive module in the unlikely event of a 2.2GB drive failure.

- Technical Description

The hardware supplied by this feature is the same as provided by feature code 3201, but includes protective packaging and installation instructions.

- Maximum number of features

One per order.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

Yes

- Limitations

Must be ordered against 7132-REP with one disk drive module per order.

- Field installable

Yes.

- Cable order

None.

FC 3401 ONE 4.5GB SSA DISK DRIVE MODULE

This feature supplies one 4.5GB SSA disk drive mounted in an auto-docking carrier. The auto-docking carrier includes a lever that facilitates safe insertion or removal of the disk drive module.

- Technical Description

The feature consists of a 4.5GB 3.5 inch form factor device mounted on a carrier. The carrier houses power conversion

circuits in addition to the disk drive. The drive characteristics are as described below:

- Formatted Capacity - 4.5GB
- Form Factor - 3.5 inch
- Average Seek Time - 8.0 msec
- Latency - 4.17 msec
- Media Data Transfer Rate - 9.59 to 12.58 MB/sec (10 bands)
- SSA Data Transfer Rate - 20 MB/sec Full Duplex
- Maximum features per model 010
Sixteen. Note limitations of total drive modules that may be installed in the subsystem.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features
3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.
- Field installable
Yes.
- Cable order
None.

FC 3404 4.5GB SSA DISK DRIVE MODULE, ONE OF FOUR (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of four 4.5GB disk drive modules in positions 1 to 4. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered four times.

- Technical Description
The feature content is the same as feature code 3401.
- Maximum features per model 010
Four. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites

- None.
- Corequisites
 - None.
- Compatibility conflicts
 - None.
- Customer setup
 - No.
- Limitations
 - The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features: 3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.
- Field installable
 - No.
- Cable order
 - None.
- Minimum order quantity
 - Four.

FC 3408 4.5GB SSA DISK DRIVE MODULE, ONE OF EIGHT (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of eight 4.5GB disk drive modules in positions 1 to 8. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered eight times. Note that for redundant power for the positions 5 to 8, feature code 3003 Power/cooling feature 3 must be ordered.

- Technical Description
 - The feature content is the same as feature code 3401.
- Maximum number of features
 - Eight. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites
 - None.
- Corequisites
 - None.
- Compatibility conflicts
 - None.
- Customer setup
 - No.
- Limitations
 - The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives

supplied by the following features: 3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable
No.
- Cable order
None.
- Minimum order quantity
Eight.

FC 3416 4.5GB SSA DISK DRIVE MODULE, ONE OF SIXTEEN (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of sixteen 4.5GB disk drive modules in positions 1 to 16. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered sixteen times.

- Technical Description
The feature content is the same as feature code 3401.
- Maximum features per model 010
Sixteen. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites
feature code3003, Power/cooling feature 3.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16.
- Field installable
No.
- Cable order
None.
- Minimum order quantity
Sixteen.

FC 3499 4.5GB SSA DISK DRIVE MODULE REPLACEMENT

This feature supplies an on-site spare 4.5GB disk drive module which a customer can use to rapidly replace a disk drive module in the unlikely event of a 4.5GB drive failure.

- Technical Description

The hardware supplied by this feature is the same as provided by feature code 3401, but includes protective packaging and installation instructions.
- Maximum number of features

One per order.
- Prerequisites

None.
- Corequisites

None.
- Compatibility conflicts

None.
- Customer setup

Yes.
- Limitations

Must be ordered against 7132-REP with one disk drive module per order.
- Field installable

Yes.
- Cable order

None.

FC 5002 0.18M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem.

- Technical Description

The cable consists of two sets of twisted pair cable.
- Maximum number of features

As required.
- Prerequisites

None.
- Corequisites

None.
- Compatibility conflicts

None.
- Customer setup

No.
- Limitations

None.
- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5006 0.60M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem. It is specifically designed for connecting drives four and five, or twelve and thirteen, on the model 500, where the 1.0 meter cable may create problems with the rear cover of the machine.

- Technical Description

The cable consists of two sets of twisted pair cable.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5010 1.0M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem, connect to an adjacent subsystem, or system adapter.

- Technical Description

The cable consists of two sets of twisted pair cable.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5025 2.5M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem, connect to an adjacent subsystem, or system adapter.

- Technical Description

The cable consists of two sets of twisted pair cable.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5050 5.0M SSA COPPER CABLE

This feature supplies a cable that can be used to link the SSA strings in a 7133 SSA Disk Subsystem to an adjacent subsystem, or to a system adapter.

- Technical Description

The cable consists of two sets of twisted pair cables.

- Maximum number of features

As required.

- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
This feature supplies a cable.

FC 5100 10.0M SSA COPPER CABLE

This feature supplies a cable that can be used to link the SSA strings in a 7133 SSA Disk Subsystem to an adjacent subsystem, or to a system adapter.

- Technical Description
The cable consists of two sets of twisted pair cable.
- Maximum number of features
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
This feature supplies a cable.

FC 5250 25.0M SSA COPPER CABLE

This feature supplies a cable that can be used to link the SSA strings in a 7133 SSA Disk Subsystem to an adjacent subsystem, or to a system adapter.

- Technical Description
The cable consists of two sets of twisted pair cable.
- Maximum number of features
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
This feature supplies a cable.

FC 5604 2.50M SSA SUBSYSTEM TO ADDITIONAL SYSTEM CABLE

This feature supplies a cable as described for feature code 5025. This feature is used to specify cables that will attach to a second system, that is, not the system associated with this machine order, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Note that this cable should only be used for connecting to adjacent drawers within a rack. It should not be used for cabling to another rack, and will not be long enough to cable from the top to the bottom of a single rack.

- Maximum number of features
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.

- Limitations

None.

- Field installable

No.

- Cable order

This feature supplies a cable.

FC 5605 5.00M SSA SUBSYSTEM TO ADDITIONAL SYSTEM CABLE

This feature supplies a cable as described for feature code 5050. This feature is used to specify cables that will attach to a second system, that is, not the system associated with this machine order, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Note that this cable can be used to connect between any subsystems or systems in a rack. It is not recommended for connecting to other racks.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

No.

- Cable order

This feature supplies a cable.

FC 5606 10.0M SSA SUBSYSTEM TO ADDITIONAL SYSTEM CABLE

This feature supplies a cable as described for feature code 5100. This feature is used to specify cables that will attach to a second system, that is, not the system associated with this machine order, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Note that this cable is recommended for connecting to systems or subsystems in other racks.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

No.

- Cable order

This feature supplies a cable.

FC 5607 25.0M SSA SUBSYSTEM TO ADDITIONAL SYSTEM CABLE

This feature supplies a cable as described for feature code 5250. This feature is used to specify cables that will attach to a second system, that is, not the system associated with this machine order, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Note that this cable is recommended for connecting to subsystems or systems in other racks.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

No.

- Cable order

This feature supplies a cable.

FC 5704 2.5M SSA SUBSYSTEM TO SUBSYSTEM CABLE

This feature supplies a cable as described for feature code 5025. This feature is used to specify cables that will attach to a second subsystem on the same SSA loop, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting subsystems in the same rack. It should not be used between racks.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5705 5.00M SSA SUBSYSTEM TO SUBSYSTEM CABLE

This feature supplies a cable as described for feature code 5050. This feature is used to specify cables that will attach to a second subsystem on the same SSA loop, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting subsystems within a rack. It is not recommended for connections across racks.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5706 10.0M SSA SUBSYSTEM TO SUBSYSTEM CABLE

This feature supplies a cable as described for feature code 5100. This feature is used to specify cables that will attach to a second subsystem on the same SSA loop, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting subsystems in different racks.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5707 25.0M SSA SUBSYSTEM TO SUBSYSTEM CABLE

This feature supplies a cable as described for feature code 5250. This feature is used to specify cables that will attach to a second subsystem on the same SSA loop, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting subsystems in different racks.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.

- Limitations

None.

- Field installable

No.

- Cable order

This feature supplies a cable.

FC 5801 0.18M SSA LOOP TO LOOP CABLE

This feature supplies a cable as described for feature code 5002. This feature is used to complete an SSA loop between drive 8 and drive 9, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems.

- Maximum number of features

One.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

No.

- Cable order

This feature supplies a cable.

FC 5803 1.0M SSA LOOP TO LOOP CABLE

This feature supplies a cable as described for feature code 5010. This feature is used to complete an SSA loop between drive 4 and drive 5, or between drive 12 and drive 13, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems.

- Maximum number of features

Two.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5904 2.50M SSA SUBSYSTEM TO ADAPTER CABLE

This feature supplies a cable as described for feature code 5025. This feature is used to specify cables that will attach to a SSA Four Port Adapter (feature code 6214) in the attaching system, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting subsystems to systems in the same rack. However, it may not be able to connect from the lowest drawer position to a system.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5905 5.00M SSA SUBSYSTEM TO ADAPTER CABLE

This feature supplies a cable as described for feature code 5050. This feature is used to specify cables that will attach to a SSA Four Port Adapter (feature code 6214) in the attaching system, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting subsystems to systems within the same rack.

- Maximum number of features

Eight.

- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5906 10.0M SSA SUBSYSTEM TO ADAPTER CABLE

This feature supplies a cable as described for feature code 5100. This feature is used to specify cables that will attach to a SSA Four Port Adapter (feature code 6214) in the attaching system, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting a subsystem to a system in a different rack.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

FC 5907 25.0M SSA SUBSYSTEM TO ADAPTER CABLE

This feature supplies a cable as described for feature code 5250. This feature is used to specify cables that will attach to a SSA Four Port Adapter (feature code 6214) in the attaching system, and is required at time of original machine order for 7133 Model 010 SSA Disk Subsystems. Each machine order must include one of the 57XX or 59XX feature codes. Note that this cable is recommended for connecting a subsystem to a system in a different rack.

- Maximum number of features
Eight.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
No.
- Cable order
This feature supplies a cable.

A.1.9 Feature Exchanges

Not available.

A.1.10 Accessories

None.

A.1.11 Customer Replacement Parts

None.

A.1.12 Machine Elements

None.

A.1.13 Supplies

None.

A.1.14 Diskettes

None.

A.2 IBM 7133 SSA Disk Subsystem Model 500

The 7133 provides power and packaging for SSA disk drive modules. The Model 500 is a standalone unit similar in height to RISC System/6000 deskside systems.

The Model 500 provides slots for up to 16 SSA disk drive modules mounted in convenient carriers. These disk drive modules are auto-docking and may be replaced by the customer. The drives connect to four internal SSA loops which can be configured as independent loops or linked via external cables.

The 7133 can be populated with three different SSA disk drive modules. The Capacities supported are 1.1GB, 2.2GB and 4.5GB. The machine includes four 2.2GB disk drives as standard. These can be replaced by eight 1.1GB disk drive modules or four 4.5GB drives at time of original order. Features are provided to add any combinations of supported drives.

The 7133 includes two auto-docking power/cooling units. These will provide redundant power and cooling to disk drive module positions 1 to 4. They will also provide non-redundant power/cooling to positions 5 to 8. An additional unit is available as a selective feature, and this will provide redundant power and cooling to all 16 disk drive module positions. These units may also be replaced by the customer.

The 7133 has selectable cable features, up to 25 meters long, for attaching to an IBM High Performance 4-Port SSA Adapter in the attaching system, for linking the internal buses, and for connecting to other 7133 subsystems.

A.2.1 Physical Specification

- Width: 210mm (8.3 inches)
- Depth: 820mm (32.3 inches)
- Height: 610mm (24.0 inches)
- Weight: 54.5Kg (120 lbs) Entry configuration
- Weight: 72.5Kg (160 lbs) Maximum configuration

A.2.2 Operating Environment

- Temperature: 16 to 32 degrees C
- Relative Humidity: 8 to 80%, non-condensing.
- Wet Bulb: 23 degrees C Calorific value: 1810 BTU/hr (maximum config)
- Electrical power : 0.645KVA (maximum config, start up)
- Electrical power : 0.480KVA (maximum config, operating)
- Electrical power : 0.420KVA (maximum config, idling)
- Capacity of Exhaust : 25 litres/second
- Noise Level : 5.4Bels (maximum config, idling)
- Noise Level : 5.6Bels (maximum config, operating)

The EMC Conformance Classification:

- USA - FCC Class A
- Germany - VDE Class A
- Japan - VCC-1

The product safety/country testing/certification:

- USA - UL
- Canada - CSA
- Germany - GS Mark (Safety, TUV)
- International - IEC

A.2.3 Limitations

Not applicable.

A.2.4 Hardware Requirements

The 7133 Model 500 attaches to the following RISC System/6000 systems:

- 7012 Models 34H, 340, 350, 360, 370, 380, 390, and 39H.
- 7013 Models 52H, 53E, 53H, 55E, 55S, 55L, 550, 56F, 560, 57F, 570, 58F, 58H, 580, 59H, 590, and 591.

The attaching system must provide an attachment to an SSA loop via feature code6214, the SSA adapter.

A.2.5 Software Requirements

The IBM 7133 SSA Disk Subsystem is supported by AIX V3.2.5 with additional PTFs, and AIX 3.2.5 device driver shipped with the SSA Four Port Adapter (feature code 6214 on the attaching system). For ease of installation these PTFs are packaged with the Device Driver on the CD-ROM shipped with the adapter. Customers without access to CD-ROM drives on their machines or network can obtain the driver and the required PTFs via FIXDIST. The device driver is available as APAR IX52018. The required PTFs, on FIXDIST, are identified as PMP3251. For alternative delivery, contact your Software Service representatives for the appropriate PTFs. The additional Version 3.2.5 PTFs (without the AIX 3.2.5 Device Driver for the adapter) are included on all AIX new V3 orders shipped after May 19, 1995 labelled AIX 3.2.5 Enhancement 5 (3250-05-00).

A.2.6 Electronic Customer Support

Not available.

A.2.7 Publications

The publications listed in Table 15 on page 141 are shipped with the product. Additional copies are available immediately.

<i>Table 15. Publications shipped with the 7133</i>	
Title	Order Number
Operator Guide	GA33-3259
Installation Guide (Model 500)	GA33-3263
Service Guide	SY33-0185
Note: Only the Ship Pack publications appear in the Sales Manual. A complete list of publications is provided in the PUBS section of HONE.	

A.2.8 Features

The following features are available for the 7133 model 500:

A.2.9 No Charge Specify Codes

The features listed in Table 16 summarize the no charge specify codes for the 7133 model 500.

<i>Table 16. No charge feature code summary</i>	
Description	Feature
Attachment System - RISC System Model 7012	7012
Attachment System - RISC System Model 7013	7013
Standard 2.2GB SSA disk drive modules, one of four	9204
100/200V 50/60Hz and 300V DC input power	9850
Language Groups - US English	9300
Language Groups - French	9703
Language Groups - German	9704
Language Groups - Spanish	9708
Language Groups - Italian	9711
Language Groups - Japanese	9714
Power cord	9800
Power cord (Chicago)	9986

The following list provides details of the no charge specify codes:

SC 9850 100/200V 50/60HZ and 300VDC input power

This feature supplies two auto-sensing power and cooling modules that will operate with input ac power between 100 and 200V, or 300V DC. This feature is mandatory for all orders.

SC 9204 Standard 2.2GB SSA disk drive modules, one of four

This feature provides 2.2GB SSA disk drive modules mounted in positions 1 to 4. They are included as standard, and may be replaced by 1.1GB or 4.5GB SSA disk drive modules at time of original order only. The specification of the disk drive modules is as described under feature code 3201. This code will be specified four times if none of the disk drive select features are ordered.

SC 7012 System attachment - RISC System/6000 model 7012

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 SSA Disk Subsystem with new RISC System/6000 7012 systems orders. This will ensure that the product will be recorded against the intended RS/6000 SP or

RS/6000 system on which it will be installed. All orders must specify either this code, feature code 9076 or feature code 7015 or feature code 7013.

SC 7013 System attachment - RISC System/6000 model 7013

This specify code is for administrative purposes only. This code must be added to all US/AP orders for the IBM 7133 SSA Disk Subsystem with new RISC System/6000 7013 systems orders. This will ensure that the product will be recorded against the intended RS/600 SP or RS/6000 system on which it will be installed. All orders must specify either this code, feature code 9076 or feature code 7015 or feature code 7012.

A.2.10 Chargeable Feature Codes

The features listed in Table 17 summarize the chargeable specify feature codes for the 7133 model 500.

<i>Table 17. Chargeable feature code summary for the model 500</i>	
Description	Feature
Remote power-on control	3001
Power/cooling feature 3	3003
Disk Drive Module Lock	3004
One 1.1GB SSA disk drive module	3101
1.1GB SSA disk drive module, one of eight (select)	3108
1.1GB disk drive module, one of sixteen (select)	3116
One 2.2GB SSA disk drive module	3201
2.2GB SSA disk drive module, one of four	3208
2.2GB SSA disk drive module, one of twelve	3216
One 4.5GB SSA disk drive module	3401
4.5GB SSA disk drive module, one of four (select)	3404
4.5GB SSA disk drive module, one of eight (select) 3408	
4.5GB SSA disk drive module, one of sixteen (select)	3416
0.18m SSA copper cable	5002
0.60m SSA copper cable	5006
1.00m SSA copper cable	5010
2.50m SSA copper cable	5025
5.00m SSA copper cable	5050
10.0m SSA copper cable	5100
25.0m SSA copper cable	5250
The following features must be ordered against the 7132-REP machine	
1.1GB SSA disk drive module replacement	3199
2.2GB SSA disk drive module replacement	3299
4.5GB SSA disk drive module replacement	3499
Power/cooling replacement unit	3099
Disk Drive Module Lock	3004

The following list provides details of the chargeable specify codes:

FC 3001 REMOTE POWER-ON CONTROL

This feature allows the attaching system(s) to control the power on/off status of the 7133 SSA Disk Subsystem.

- Technical description

The feature will control the power on /off status of the power supplies within the 7133 SSA Disk Subsystem under control of up to two systems. When two systems are connected, the product

will power on if requested by either host, and will power off when requested by both.

- Maximum features per model 500
One.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
None.

FC 3003 POWER/COOLING FEATURE 3

This feature provides redundant power and cooling to disk drive positions 5 to 16.

- Technical description
The 7133 SSA Subsystem is supplied with two auto-docking combined power/cooling units. These supply redundant power/cooling to disk drive positions 1 to 4, and non-redundant power/cooling to positions 5 to 8. This feature is mandatory for disk drive positions 9 to 16. This feature provides a third auto-docking unit which supplies redundant power/cooling to the complete subsystem.
- Maximum features per model 500
One.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.

- Field installable
Yes.
- Cable order
None.

FC 3004 DISK DRIVE MODULE LOCK

This feature provides a lock which can be used to control removal of the disk drive modules. The first order of this feature will provide three keys and a lock. Subsequent orders in the same total order will provide only a lock.

- Technical description
Disk drive modules are fitted with a lever which ensures correct connection of the module on insertion, and controlled removal when required. The lock supplied by this feature fits on the front of the module and inhibits the use of the lever. There is only one version of the lock. This feature prevents inadvertent removal of the module. It cannot be considered as protection against malicious removal.
- Maximum features per model 500
Sixteen.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
None.

FC 3099 POWER/COOLING REPLACEMENT UNIT

This feature supplies an on-site spare power/cooling unit which a customer can use to rapidly replace a power/cooling unit in the unlikely event of a failure.

- Technical description
The hardware supplied by this feature is the same as provided by feature code 3003, but includes protective packaging, and installation instructions.
- Maximum features per model 500

- As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
Yes
- Limitations
Must be ordered against 7132-REP with one power/cooling unit per order.
- Field installable
Yes.
- Cable order
None.

FC 3101 ONE 1.1GB SSA DISK DRIVE

This feature supplies one 1.1GB SSA disk drive mounted in an auto-docking carrier. The auto-docking module includes a lever that facilitates safe insertion or removal of the disk drive module.

- Technical description
The feature consists of a 1.1GB 3.5 inch form factor device mounted on a carrier. The carrier houses power conversion circuits in addition to the disk drive. The drive characteristics are as described below:
 - Formatted Capacity - 1.1GB
 - Form Factor - 3.5inch
 - Average Seek Time - 6.9 msec
 - Latency - 4.17 msec
 - Media Data Transfer Rate - 9.59 to 12.58 MB/sec (10 bands)
 - SSA Data Transfer Rate - 20 MB/sec Full Duplex
- Maximum features per model 500
Sixteen. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.

- Customer setup
No.
- Limitations
The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features
3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.
- Field installable
Yes.
- Cable order
None.
- Maximum order quantity
Sixteen

FC 3108 1.1GB SSA DISK DRIVE MODULE, ONE OF EIGHT (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of eight 1.1GB drives in positions 1 to 8. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered eight times. Note that for redundant power for the positions 5 to 8, feature code 3003 Power/cooling feature 3 must be ordered.

- Technical description
The feature content is the same as feature code 3101.
- Maximum features per model 500
Eight. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features
3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.
- Field installable
No.
- Cable order

None.

- Minimum order quantity

Eight.

FC 3116 1.1GB SSA DISK DRIVE MODULES, ONE OF SIXTEEN (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive normally included with the machine and installation of sixteen 1.1 GB drives in positions 1 to 16. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered sixteen times.

- Technical description

The feature content is the same as feature code 3101.

- Maximum features per model 500

Sixteen. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

feature code 3003, Power/cooling feature 3.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Sixteen.

FC 3199 1.1GB SSA DISK DRIVE MODULE REPLACEMENT

This feature supplies an on-site spare 1.1GB disk drive module which a customer can use to rapidly replace a disk drive module in the unlikely event of a 1.1GB drive failure.

- Technical description

The hardware supplied by this feature is the same as provided by feature code 3101, but includes protective packaging, and installation instructions.

- Maximum number of features
One per order.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
Yes
- Limitations
Must be ordered against 7132-REP with one disk drive module per order.
- Field installable
Yes.
- Cable order
None.

FC 3201 ONE 2.2GB SSA DISK DRIVE MODULE

This feature supplies one 2.2GB SSA disk drive mounted in an auto-docking carrier. The auto-docking carrier includes a lever that facilitates safe insertion or removal of the disk drive module.

- Technical description
The feature consists of a 2.2GB 3.5 inch form factor device mounted on a carrier. The carrier houses power conversion circuits in addition to the disk drive. The drive characteristics are as described below:
 - Formatted Capacity - 2.2GB
 - Form Factor - 3.5inch
 - Average Seek Time - 7.5 msec
 - Latency - 4.17 msec
 - Media Data Transfer Rate - 9.59 to 12.58 MB/sec (10 bands)
 - SSA Data Transfer Rate - 20 MB/sec Full Duplex
- Maximum features per model 500
Sixteen. Note limitations of total drive modules that may be installed in the subsystem.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

Yes.

- Cable order

None.

FC 3208 2.2GB SSA DISK DRIVE MODULES, ONE OF FOUR

Ordering this feature results in the installation of four 2.2GB drives in positions 5 to 8. This feature is available at time of first order only. Note that for redundant power for the positions 5 to 8, feature code 3003 Power/cooling feature 3 must be ordered.

- Technical description

The feature content is the same as feature code 3201.

- Maximum features per model 500

Four. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Four.

FC 3216 2.2GB SSA DISK DRIVE MODULES, ONE OF TWELVE

Ordering this feature results in the installation of twelve 2.2GB drives in positions 5 to 16. This feature is available at time of first order only and must be ordered twelve times.

- Technical description

The feature content is the same as feature code 3201.

- Maximum features per model 500

Twelve. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

feature code 3003, Power/cooling feature 3.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Twelve.

FC 3299 2.2GB SSA DISK DRIVE MODULE REPLACEMENT

This feature supplies an on-site spare 2.2GB disk drive module which a customer can use to rapidly replace a disk drive module in the unlikely event of a 2.2GB drive failure.

- Technical description

The hardware supplied by this feature is the same as provided by feature code 3201, but includes protective packaging and installation instructions.

- Maximum number of features

One per order.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

Yes

- Limitations

Must be ordered against 7132-REP with one disk drive module per order.

- Field installable

Yes.

- Cable order

None.

FC 3401 ONE 4.5GB SSA DISK DRIVE MODULE

This feature supplies one 4.5GB SSA disk drive mounted in an auto-docking carrier. The auto-docking carrier includes a lever that facilitates safe insertion or removal of the disk drive module.

- Technical description

The feature consists of a 4.5GB 3.5 inch form factor device mounted on a carrier. The carrier houses power conversion circuits in addition to the disk drive. The drive characteristics are as described below:

- Formatted Capacity - 4.5GB
- Form Factor - 3.5inch
- Average Seek Time - 8.0 msec
- Latency - 4.17 msec
- Media Data Transfer Rate - 9.59 to 12.58 MB/sec (10 bands)
- SSA Data Transfer Rate - 20 MB/sec Full Duplex

- Maximum features per model 500

Sixteen. Note limitations of total drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

Yes.

- Cable order

None.

FC 3404 4.5GB SSA DISK DRIVE MODULE, ONE OF FOUR (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of four 4.5GB disk drive modules in positions 1 to 4. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and be ordered four times.

- Technical description

The feature content is the same as feature code 3401.

- Maximum features per model 500

Four. Note limitations of total disk drive modules that may be installed in the subsystem.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features

3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Four.

FC 3408 4.5GB SSA DISK DRIVE MODULE, ONE OF EIGHT (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of eight 4.5GB disk drive modules in positions 1 to 8. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered eight times. Note that for redundant power for the positions 5 to 8, feature code 3003 Power/cooling feature 3 must be ordered.

- Technical description
The feature content is the same as feature code 3401.
- Maximum number of features
Eight. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16. This total will include drives supplied by the following features
3101, 3108, 3116, 3201, 3208, 3216, 3401, 3404, 3408, 3416 and 9204.
- Field installable
No.
- Cable order
None.
- Minimum order quantity
Eight.

FC 3416 4.5GB SSA DISK DRIVE MODULE, ONE OF SIXTEEN (SELECT)

Ordering this feature results in the removal of the four 2.2GB SSA disk drive modules normally included with the machine and installation of sixteen 4.5GB disk drive modules in positions 1 to 16. The feature is priced to provide credit for the removal of the 2.2GB drives. This feature is available at time of first order only and must be ordered sixteen times.

- Technical description
The feature content is the same as feature code 3401.
- Maximum features per model 500
Sixteen. Note limitations of total disk drive modules that may be installed in the subsystem.
- Prerequisites
feature code 3003, Power/cooling feature 3.
- Corequisites
None.
- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

The maximum number of disk drive modules that can be installed in a 7133 SSA Disk Subsystem is 16.

- Field installable

No.

- Cable order

None.

- Minimum order quantity

Sixteen.

FC 3499 4.5GB SSA DISK DRIVE MODULE REPLACEMENT

This feature supplies an on-site spare 4.5GB disk drive module which a customer can use to rapidly replace a disk drive module in the unlikely event of a 4.5GB drive failure.

- Technical description

The hardware supplied by this feature is the same as provided by feature code 3401, but includes protective packaging and installation instructions.

- Maximum number of features

One per order.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

Yes

- Limitations

Must be ordered against 7132-REP with one disk drive module per order.

- Field installable

Yes.

- Cable order

None.

FC 5002 0.18M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem.

- Technical description
The cable consists of two sets of twisted pair cable.
- Maximum number of features
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
This feature supplies a cable.

FC 5006 0.60M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem. It is specifically designed for connecting drives four and five, or twelve and thirteen, on the model 500, where the 1.0meter cable may create problems with the rear cover of the machine.

- Technical description
The cable consists of two sets of twisted pair cable.
- Maximum number of features
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5010 1.00M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem, connect to an adjacent subsystem, or system adapter.

- Technical description

The cable consists of two sets of twisted pair cable.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5025 2.50M SSA COPPER CABLE

This feature supplies a cable that can be used to link the internal SSA strings in a 7133 SSA Disk Subsystem, connect to an adjacent subsystem, or system adapter.

- Technical description

The cable consists of two sets of twisted pair cable.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
This feature supplies a cable.

FC 5050 5.00M SSA COPPER CABLE

This feature supplies a cable that can be used to link the SSA strings in a 7133 SSA Disk Subsystem to an adjacent subsystem, or to a system adapter.

- Technical description
The cable consists of two sets of twisted pair cables.
- Maximum number of features
As required.
- Prerequisites
None.
- Corequisites
None.
- Compatibility conflicts
None.
- Customer setup
No.
- Limitations
None.
- Field installable
Yes.
- Cable order
This feature supplies a cable.

FC 5100 10.0M SSA COPPER CABLE

This feature supplies a cable that can be used to link the SSA strings in a 7133 SSA Disk Subsystem to an adjacent subsystem, or to a system adapter.

- Technical description
The cable consists of two sets of twisted pair cable.
- Maximum number of features
As required.
- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

This feature supplies a cable.

FC 5250 25.0M SSA COPPER CABLE

This feature supplies a cable that can be used to link the SSA strings in a 7133 SSA Disk Subsystem to an adjacent subsystem, or to a system adapter.

- Technical description

The cable consists of two sets of twisted pair cable.

- Maximum number of features

As required.

- Prerequisites

None.

- Corequisites

None.

- Compatibility conflicts

None.

- Customer setup

No.

- Limitations

None.

- Field installable

Yes.

- Cable order

This feature supplies a cable.

A.2.11 Feature Exchanges

Not available.

A.2.12 Accessories

None.

A.2.13 Customer Replacement Parts

None.

A.2.14 Machine Elements

None.

A.2.15 Supplies

None.

A.2.16 Diskettes

None.

Appendix B. 7133 Configuration Planning Charts

This chapter provides configuration charts that can be used to assist in designing and documenting 7133 SSA configurations.

B.1 Planning Chart

Copy this page as many times as required and then number each adapter and subsystem uniquely. Next, use the table following to identify the types of disk in each slot, the systems each adapter is in and the cabling connections.

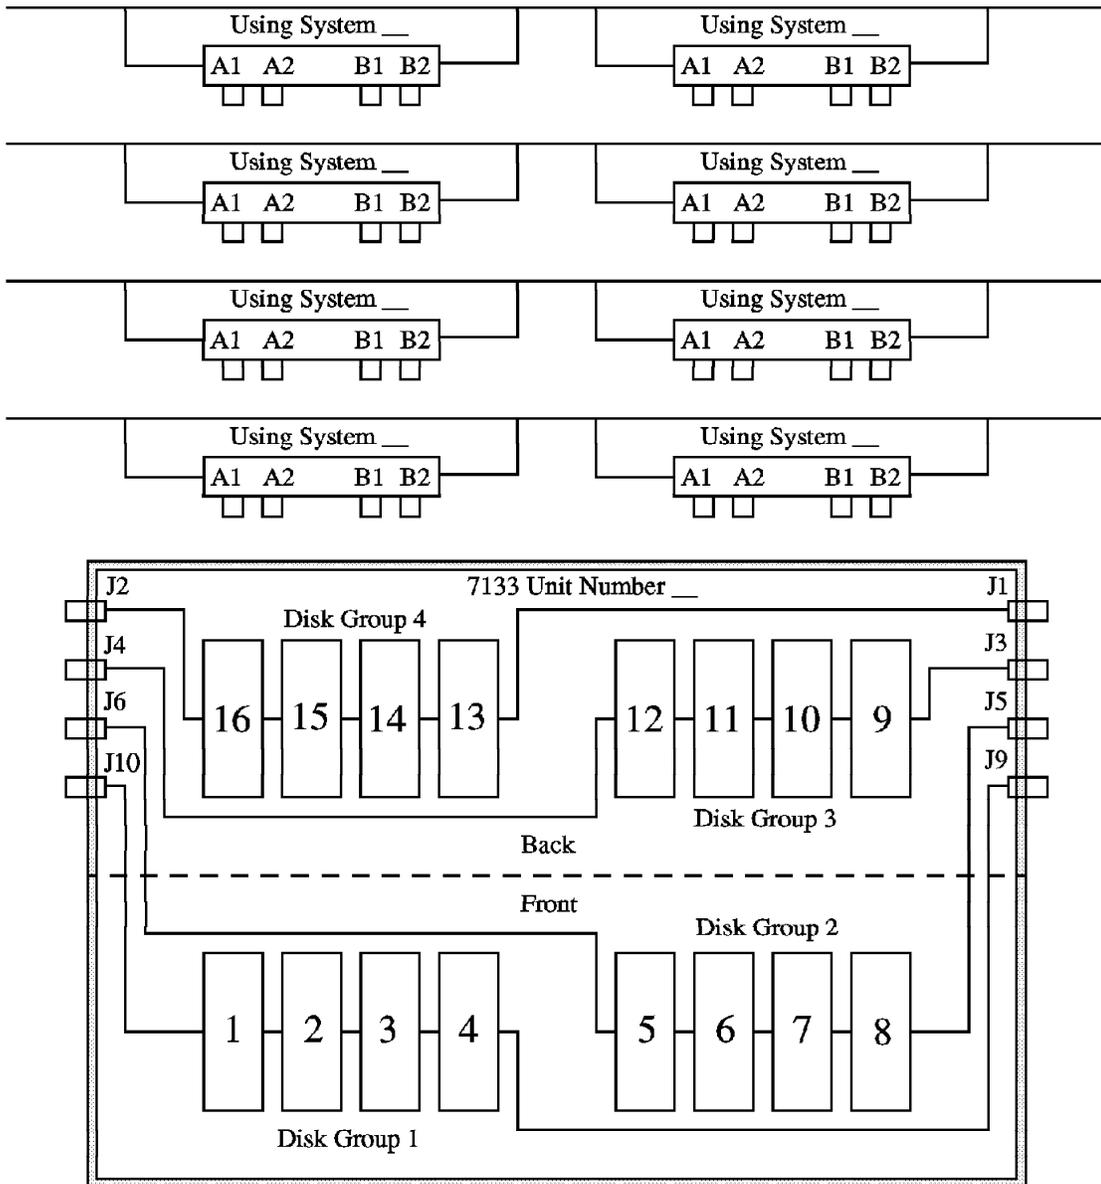


Figure 48. Configuration Planning Chart Page ___

B.2 Configuration Planning Tables

Use this table to indicate which disks are in each slot, which host systems adapters belong to and the cabling connections.

B.2.1 Adapter Locations and Connections

Use this table to record the location of adapters and the port connections from the adapters.

Table 18. Adapter Location and Connections. The SSA four-port adapter host system location and connections.

Adapter Number	Host System Name	Connections			
		A1	A2	B1	B2
1		___on___	___on___	___on___	___on___
2		___on___	___on___	___on___	___on___
3		___on___	___on___	___on___	___on___
4		___on___	___on___	___on___	___on___
5		___on___	___on___	___on___	___on___
6		___on___	___on___	___on___	___on___
7		___on___	___on___	___on___	___on___
8		___on___	___on___	___on___	___on___

Note:

1. Maximum of three adapters per system
2. Maximum of two adapters per loop
3. Maximum of eight hosts per subsystem (four loops, two hosts per loop)

B.2.2 7133 Subsystem Configuration and Connections

Use this table to record the 7133 subsystem configuration and connections.

Subsystem Number ____				
Disk Slot	Disk Type	Connect From	Connect To	Cable Type
1		J10	____ on ____	
2				
3		J9	____ on ____	
4				
5		J6	____ on ____	
6				
7		J5	____ on ____	
8				
9		J3	____ on ____	
10				
11		J4	____ on ____	
12				
13		J1	____ on ____	
14				
15		J2	____ on ____	
16				
<p>Note:</p> <ol style="list-style-type: none"> 1. Maximum of 48 disks on a loop 2. Maximum of 25 meters between connections 3. Use the tables on 2.1.2.4, "7133 Hardware Schematics and Connection Rules" on page 24 to select cables 				

B.2.3 Topology Configuration Chart

Use this chart to sketch out the prospective layout of the configuration. Links can be drawn between the various elements. If more subsystems are needed, use multiple copies of this page.

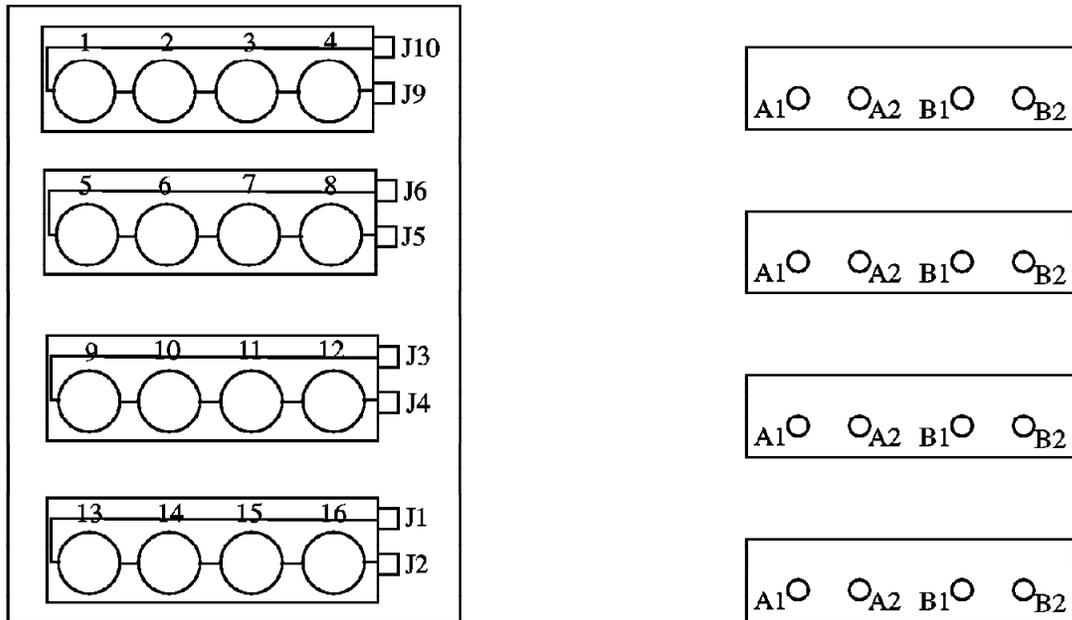


Figure 49. Topology Planning Chart Page ___

Remember to document the configuration.

List of Abbreviations

APAR	Authorized Program Analysis Report	NVS	Non-Volatile Storage
ASIC	Application Specific Integrated Circuits	ODM	Object Data Manager
DSV	Digital Sum Variation	PMP	Preventative Maintenance Procedures
DRAM	Dynamic Random Access Memory	POST	Power-On Self Test
EC	Engineering Change	PTF	Program Temporary Fix
EMI	ElectroMagnetic Interference	RAID	Redundant Array of Independent Disks
ERP	Error Recovery Procedures	ROS	Read-Only Storage
FCS	Fiber Channel Switch	RR	Receiver Ready
FDDI	Fiber Distributed Data Interface	SCSI	Small Computer System Interface
FRU	Field Replaceable Unit	SRN	System Reference Number
IBM	International Business Machines Corporation	SSA	Serial Storage Architecture
ITSO	International Technical Support Organization	TPS	Transactions Per Second
LPP	Licensed Program product	VGDA	Volume Group Descriptor Area
MAP	Maintenance Analysis Procedures	VGSA	Volume Group Status Area
		UIG	User's Industry Group

Index

Special Characters

????? 84, 93
***** 84

Numerics

7133 SSA Disk Subsystem
 activation 15, 18
 basic configuration 40
 cables 25
 connection rules 24
 deactivation 15, 18
 disk drive module physical layout 22
 four-port adapter physical layout 23
 hardware schematics 24
 internal disk numbering 27
 loop rules 29
 model 010 physical layout 15
 model 500 physical layout 17
 operating system support 14
 power and cooling configuration rules 15
 problems 97
 product overviews 13
 technical details 14
8B/10B coding 6
9333 Serial Disks
 architecture 2
 use 3

A

A_Quota 11
Abbreviations 165
ABORT 7, 9
ACK 7, 8
Acronyms 165
Adapter 86
Adapter ownership 73
Adapter Port 84
adapter_a 47
adapter_b 47
ASIC 2
assign 47
Auto-docking 14
Availability and performance related scenarios 72
Availability-related scenarios 60
Asynchronous alert 10

B

B_Quota 11
Bandwidth 49, 72, 75
Block size 101

C

Cables 25
 for complete system orders 25
 for stand-alone models 25
 link problems 95
Certification time 91
cfgmgr 40, 44
Changing the queue depth 107
Channel 8
chdev 47, 107
Check light 22
chpv 107
Clock deviation 1
Comma codes 6
Communications errors 97
Configuration 39
Configuration planning charts and tables 161
Connection Rules 24
 connecting groups between subsystems 27
 connecting groups inside a subsystem 26
 connecting groups to adapters 28
 creating loops 29
Connectivity 1
Cost 72, 73
Creating loops 29
Crosstalk 1
Cut-through routing 9
Cyclic network 5
Cyclic path 5

D

Data location 53, 61, 74
Data path 51, 73
DC bias 6
Decoding port and SSA addresses 92
Dedicated connection 4
Device Driver Level 35
diag 41, 44, 77
Diagnostics 31
Differential pairs 2
Digital sum variation 6
DIS 7
Disk position 73
Displaying link status 82
Displaying power status 82
Displaying queue depth 107
Documentation required 33
Double failure 74
Double failures 73, 74
Dual port 4

E

EC Level 35
Electromagnetic interference 1
Error log 94
Error rate 10

F

Failed 84
Fallback 72
Fan-and-power check light 16, 19
Fiber Channel 2
File size 100
FLAG 7
Form factor 4
Formatting times 89
Frame 4
Frames
 address field 8
 control byte 7
 CRC field 8
 data field 8
 format 8
 sequence number 8
 types 7
FRU list 94
Full duplex 2

G

Good 84
Ground bounce 1

H

HACMP 60, 72
Hardware scenarios 49
Hardware schematics 24
hdisk 30, 40, 46
Higher-level requirements 54
History of SSA 1
Hot plugging 4, 10, 42

I

Identify mode 81
IEEE P1394 2
Installation and Configuration
 adding disks 42
 basic configuration steps 39
 changing the path to a disk 47
 checking adapter levels 34
 checking software levels 33
 configuration verification 41
 documentation required 33
 hardware required 35
 installation steps 38
 listing devices created 40
 moving disks 46

Installation and Configuration (*continued*)

 prerequisites 38
 removing disks 44
 software required 33

L

Listing physical drives 79
Load balancing 52
Loadable Microcode Level 35
Locating devices without the Service Aids 92
Locating disks 81
Logical volumes 54, 61
Loop 2, 5
Loop Configurations
 availability-related 60
 basic eight disk loop 43
 multiple hosts 72
 one host with more than one adapter 66
 one host with one adapter 63
 performance-related 49
 random data access profile 57
 sequential data access profile 54
 two basic four disk two loops 44
 two hosts with more than two adapters 72
 two hosts with two adapters 71
lsattr 107
lscfg 34
lsdev 40, 44
lslpp 33

M

Manufacturer 35
Media errors 97
migratepv 44
Mirroring 61, 63, 74, 75
Model 010
 back connectors 17
 hardware schematics 24
 installation hardware requirements 35
 lights and power location 15
 parts location 16
 reference information 109
Model 500
 back connectors 20
 hardware schematics 24
 installation hardware requirements 36
 lights and power location 17
 parts location 19
 reference information 139
Multiple adapters 52, 60, 73, 75
Multiple disk types 54, 62, 74
Multiple loops 61, 75
Multiple systems access 60, 74

N

NUL 7, 9

O

Object data manager 46

P

Parallel mirroring 74

Part Number 34

Path 8

pdisk 30, 40, 46

Percentage 94

Performance and availability related scenarios 72

Performance-related scenarios 49

Performance/Availability Investigation

bandwidth 99

data path variations 102

effect of data location 106

loop throughput 101

single disk throughput 100

tuning queue depth 107

varying disk type 106

varying the number of adapters 105

Philips Inter-Integrated-Circuit bus 2

Physical 83, 86

Port 5, 86

POST 94

Power 84

Power and cooling 14, 15

Power card light 16, 19

Power light 16, 19, 22

Power switch 15, 18

Problem 94

Problem Determination

adapter problems 96

disk problems 96

hardware problems 97

link problems 95

problem determination procedures 94

problem solving procedures 94

Service Aid errors 92

software problems 95

Product overviews 13

Product reference information 109

PTF requirements 33

Q

Queue depth 107

Quorum 62

R

RAID/SSA Configuration Tool

Read performance 74

Read/write ratio 100

Ready light 22

reducevg 44

Remote-power-on control unit 16, 19

Reserved 84

rmdev 44, 46

ROS level and ID 35

Routing 51

RR 7, 8

RS232 2

Run-length limit 6

S

SAT 7, 11

SAT' 7, 11

Satisfaction 51

SCSI Adapters

performance 49

Seek time 100

Self-clocking 3

Sequential mirroring 74

Sequential/random I/O 101

Serial Bus

architecture 2

examples 2

Serial link advantages 1

Serial links 1

Serial Number 34

Serial Storage Architecture

characteristics 4

coding 5

component performance 50

cut-through routing 9

data rate 3

distance 3

error rate 10

error recovery procedures 10

flow control 8

flow of ACKs and RRs 9

frames 7

loops 5

node types 4

overview 3

physical layer 3

privileged messages 10

SAT algorithm 11

SCSI comparison 49

SCSI-2 mapping 3

spatial reuse 11

special flags 7

strings 4

switches 5

topology 4

transport layer 3

Serial# 84, 86

Service Aids

certify disk 90

configuration verification 85

- Service Aids (*continued*)
 - errors 92
 - format disk 87
 - link verification 41, 82
 - locating devices without the Service Aids 92
 - resolving problems 94
 - set or reset identify 81
 - set or reset service mode 81
 - set service mode 44, 79
 - starting 77
 - using 79
- Service mode 81
- Single point of failure 61, 63, 66, 71
- Single port 4
- Slicing point 6
- Software requirements 33
- Software support 30
- Spatial reuse 11
- SRN 92, 94
- ssa 33, 38
- SSA Disk Drive Modules
 - adding disks 42
 - certification 42
 - changing the path to a disk 47
 - check light state 22
 - determining location code 80, 85
 - determining pdisk/hdisk relationship 85
 - determining SSA serial number 80, 85
 - indicator lights 22
 - internal disk numbering 27
 - listing logical drive status 85
 - listing logical drives 85
 - listing physical drive status 82, 84, 86
 - listing physical drives 79
 - locating physical drives 81
 - logical disk device driver 30
 - loop rules 29
 - module locations 24
 - moving disks 46
 - part numbers 97
 - physical disk device driver 30
 - problems 96
 - ready light state 22
 - removing disks 44
- SSA Four-Port Adapter
 - connectors 24
 - determining location code 83
 - determining resource name 83
 - device driver 30
 - indicator lights 24
 - listing adapters 83
 - primary/secondary adapter 47
 - problems 96
- SSA router 30
- SSA Software Support
 - configuration methods 30
 - device drivers 30
 - diagnostics 31

- SSA Software Support (*continued*)
 - problems 95
 - SSA router 30
 - SSA-PH1 3
 - SSA-S2P 3
 - SSA-TL1 3
 - SSA_Addr 86
 - SSA01 92
 - SSA02 93
 - SSA03 93
 - ssadisk 33, 38
 - ssar 30
 - Status 84, 86
 - Status bar 90, 91
 - String 4
 - Subsystem check light 15, 18
 - Subsystem power light 15, 18
 - Switch 4, 5
 - SYNC 7
 - Synchronizing disk drive rotation 7

T

- Takeover 72
- Technology overview 1
- Third-party SSA products 31
- Transactions 75

U

- Ultrastar drives 13

V

- VDSA 62
- VGDA 62
- Volume groups 54, 61

W

- Wormhole routing 9, 51
- Write performance 74

ITSO Redbook Evaluation

**International Technical Support Organization
A Practical Guide to
Serial Storage Architecture for AIX
July 1996**

Publication No. SG24-4599-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please fill out this questionnaire and return it using one of the following methods:**

- Mail it to the address on the back (postage paid in U.S. only)
- Give it to an IBM marketing representative for mailing
- Fax it to: Your International Access Code + 1 914 432 8246
- Send a note to REDBOOK@VNET.IBM.COM

**Please rate on a scale of 1 to 5 the subjects below.
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)**

Overall Satisfaction	_____		
Organization of the book	_____	Grammar/punctuation/spelling	_____
Accuracy of the information	_____	Ease of reading and understanding	_____
Relevance of the information	_____	Ease of finding information	_____
Completeness of the information	_____	Level of technical detail	_____
Value of illustrations	_____	Print quality	_____

Please answer the following questions:

- a) Are you an employee of IBM or its subsidiaries: Yes____ No____
- b) Do you work in the USA? Yes____ No____
- c) Was this redbook published in time for your needs? Yes____ No____
- d) Did this redbook meet your needs? Yes____ No____

If no, please explain:

What other topics would you like to see in this redbook?

What other redbooks would you like to see published?

Comments/Suggestions: (THANK YOU FOR YOUR FEEDBACK!)

Name

Address

Company or Organization

Phone No.



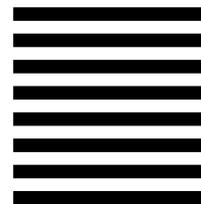
Fold and Tape

Please do not staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES



BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM International Technical Support Organization
Department JN9B, Building 045
Internal Zip 2834
11400 BURNET ROAD
AUSTIN TX
USA 78758-3493



Fold and Tape

Please do not staple

Fold and Tape



Printed in U.S.A.

SG24-4599-00

