Exploiting SSA Disk Subsystems in Sun Solaris Platform Environments

Dave McAuley, Ron Case, Gareth Donald, Jose Munoz

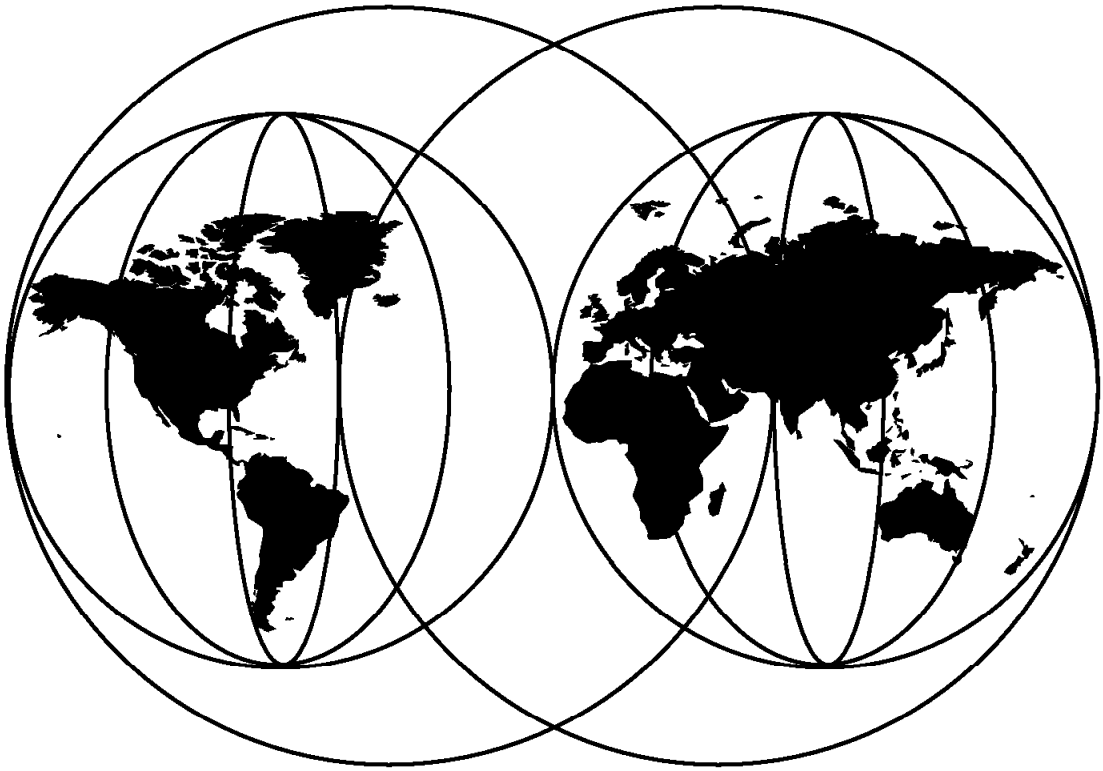

International Technical Support Organization

http://www.redbooks.ibm.com

SG24-5083-00

International Technical Support Organization

SG24-5083-00

Exploiting SSA Disk Subsystems in Sun Solaris Platform Environments

August 1998

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix A, “Special Notices” on page 197.

First Edition (August 1998)

This edition applies to the IBM SSA Interface Controller for Sun SBus and the 7190 SCSI Host to SSA Loop Attachment Model 100 for use with SCSI host servers and IBM SSA disk subsystems.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

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.

Contents

Preface

This redbook demonstrates how to install and use IBM SSA disk subsystems in non-IBM platform environments using SCSI-SSA ″bridging technology″ and native SSA adapter card technology. This bridging technology and native-SSA-attach capability enables SUN, HP, and DEC operating systems to connect to SSA disk subsystems using either native bus attachment or conversion box technology.

This redbook also provides a broad understanding of the IBM 7190 SCSI Host to SSA Loop Attachment and the IBM SSA Interface Controller for Sun SBus adapters. It will help the systems administrator know when to use these devices in the SSA configuration. These devices allow the attachment of SSA disk devices to Sun servers. This book describes the technologies, and shows the processes involved in planning and designing your configuration, as well as how to install, configure, migrate and maintain the storage network.

The Team That Wrote This Redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization San Jose Center.

Dave McAuley is the project leader for open systems disk storage products at the International Technical Support Organization, San Jose Center. He has written nine ITSO redbooks and several international conference papers, among them a paper on image processing and pattern recognition in IBM′s *Journal of Research and Development*. Dave teaches IBM classes worldwide on all areas of IBM disk storage products. Before joining the ITSO in 1994, he worked in the product support center in Scotland as a large systems storage specialist. Dave has worked at IBM for over 19 years, with international career assignments in manufacturing, design and development, and marketing. His background in computer science spans a varied portfolio including supercomputing application development, process control and image analysis systems design and microcode development, automation and robotics systems project management, and large systems marketing and technical support.

Ron Case is a Software Engineer with IBM Storage Systems Division in San Jose. As a member of the SSA Open Systems Product Field Engineering team, he provides SSA engineering and operating system expertise to IBM support staff and customers worldwide. He has eight years of systems engineering and experience spanning Windows NT, Macintosh, OpenVMS, and multiple UNIX operating systems in networked environments. He holds a BS degree in Business Administration/Computer Applications and Systems from California State University, Fresno.

Gareth Donald is the Technical manager of a leading storage distributor. He has 5 years of experience in the storage arena. He has a diploma in information technology and is a Microsoft Certified Systems Engineer. His areas of expertise are storage and backup for the midrange arena including IBM platforms such as RS/6000 and AS/400, as well as open platforms such as Solaris, HP, and DEC. He has considerable experience implementing backup and storage solutions on NT platforms. He resides in Johannesburg, South Africa where he concentrates on providing SSA storage solutions and backup solutions for large companies.

Jose Munoz is a Technical Support Manager in Bogota, Colombia. He has eight years of experience in Information Technology and telecommunications. He holds a degree in Systems and Electronics Engineering from Catholic University in Bogota, Colombia and is certified as an advanced systems administrator (ASA). His areas of expertise include UNIX, networking, and storage environments. Jose has written technical articles and bulletins on UNIX and various telecommunications topics.

This book was last updated by Mark Blunden, ITSO-San Jose.

Thanks to the following people for their invaluable contributions to this project:

John Aschoff
IBM SSD, Tucson

Maggie Cutler
ITSO

Steve Garrett
IBM SSD, San Jose

Kailash Jayaswal
IBM SSD, San Jose

Patty Krohn
ITSO

Stuart Marshall
IBM SSD, San Jose

Ling Pong
IBM ATS, San Jose

David Sacks
IBM US, Chicago

Brandon Wong
IBM SSD, Tucson

Comments Welcome

Your comments are important to us!

We want our redbooks to be as helpful as possible. Please send us your comments about this or other redbooks in one of the following ways:

- Fax the evaluation form found in "ITSO Redbook Evaluation" on page 207 to the fax number shown on the form.
- Use the electronic evaluation form found on the Redbooks Web sites:

  For Internet users `http://www.redbooks.ibm.com/`
  For IBM Intranet users `http://w3.itso.ibm.com/`

- Send us a note at the following address:

  `redbook@us.ibm.com`

Chapter 1. Overview

The IBM 7190-100 SCSI Host to SSA Loop Attachment and the IBM SSA Interface Controller for Sun SBus introduce the power and flexibility of Serial Storage Architecture (SSA) to non-IBM platforms. Although the focus of this book is the implementation of these products on Sun platforms running the Solaris operating system, many of the principles we present are independent of operating system. Thus this book is a useful reference for all administrators involved with IBM SSA technology. This chapter provides an overview of SSA, briefly describes the products available for open platforms, and explains the benefits the products have to offer.

1.1 Serial Storage Architecture

SSA is a high speed serial interface that IBM originally developed as a proprietary interface. The interface was standardized in 1994 under the control of the ANSI X3T10.1 committee. SSA has been designed to overcome the bottlenecks of existing SCSI bus architectures.

Figure 1 shows SSA in a loop topology

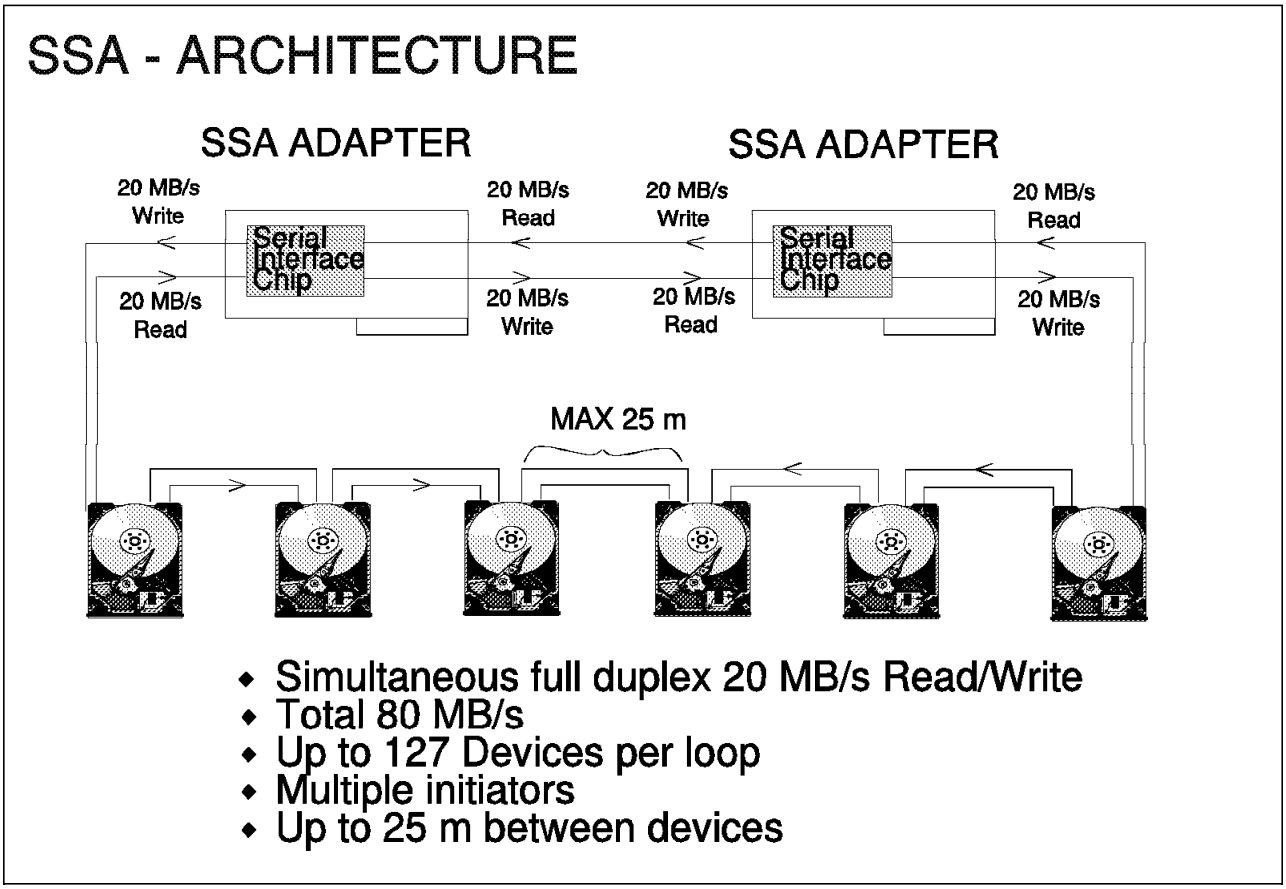

Figure 1. SSA Architecture

SSA uses a port design that enables inbound conversations and outbound conversations simultaneously. An SSA Serial Interface Chip (SIC) on an adapter that is configured into a loop has a bandwidth of 80 MB/s. There may be several

hosts on a loop with the possibility of multiple hosts in a single server, thereby enabling a server to connect to several loops. SSA connects to its transport medium by means of nine-way microminiature D connectors.

SSA uses shielded twisted pair (STP) cabling as a transport medium. There are five conductors in the cable. Four are used for signal transmission and one is used as a shield. The twists in the cable eliminate crosstalk and the shielding eliminates electromagnetic interference (EMI). Thus, SSA can function reliably with its nodes placed 25 m apart. IBM offers a fiber optic extender that can extend this limit to 2400 m. The fiber optic extenders are supplied in pairs to complete the loop.

1.1.1 Functions

Because SSA allows SCSI-2 mapping, all functions associated with initiators, targets, and logical units are translatable. Therefore, SSA can use the same command descriptor blocks, status codes, command queuing, and all other aspects of current SCSI systems.

Web configuration is built into the standard. All initiators perform a Web configuration process by which a configuration table is created. One of the initiators is selected as a Web Master. The Web master issues special configuration messages and coordinates error recovery of the Web. Unlike SCSI and fiber channel arbitrated loop (FC-AL), SSA does not use bus phases. SSA does not use arbitration to contend for the bus. Instead, it makes use of frame multiplexing. All data is transported as frames.

SSA can be set up in three different topologies:

- String topology

  A string is a simple linear network of two or more nodes. The ports at either end can be single-port nodes,while the others are dual port nodes

- Loop topology

  Loop topology is by far the most popular and most highly recommended topology. A loop is a cyclic network that contains only dual port nodes. Loops provide two data paths between nodes. The resulting benefit is higher bandwidth and higher reliability. Single node failures and single link breaks do not halt communication.

- Switch topology

  Switch topology requires additional hardware and is the most complicated to set up and maintain. Its benefit is that it is the most fault tolerant topology of the three and allows a virtually unlimited network configuration.

1.1.2 Advantages

1.1.2.1 High Performance

Only the nodes on either side of a virtual link are involved in transferring frames. A port can send a frame on its line out while it receives frames on its line in, thus enabling other nodes to initiate or process I/O requests simultaneously. This concept is known as *spatial reuse* and is very different from bus architectures that require the dedication of the entire bus for the duration of each data transfer. Fairness algorithms are used to divide bandwidth equally when there is contention for the same virtual loop, preventing a single device

from hogging a virtual link. Spindle synchronization is integrated into SSA without the requirement of additional cables. SSA bandwidth is four times that of SCSI-2 and twice that of SCSI-3. Because multiple devices can process transfers, SSA uses its bandwidth much more efficiently than SCSI. SSA overhead is approximately 6%, effectively allowing SSA to run at a 94% efficiency rate (94% of the data that SSA handles is customer data).

1.1.2.2 High availability, fault tolerance, and hot plugging

The loop and switch topologies do not have a single point of failure. As displayed in Figure 1 on page 1, the loop topology has alternate paths to each node and eliminates a single point of failure. It can connect from 2 to 128 dual port nodes. A single break in the loop does not bring down the network. Frames are automatically rerouted across the available link. SSA devices are ″hot pluggable″ by design, so nodes can be inserted into the loop without disruption. Redundant copies of data stored on SSA can be placed up to 2400 m away. SSA used in RAID systems makes up for overhead with its high speed and user data efficiency.

1.1.2.3 Greater connectivity and ease of use

SSA allows the connection of 2 to 129 nodes in a string, 2 to 128 devices in a loop, and virtually unlimited nodes on a switch. Compare this with the only 16 devices for SCSI. SSA supports multiple initiators. Manual ID selection and termination are not required, as these are automated functions of SSA and minimize the possibility of configuration errors.

1.1.2.4 Reliability

SSA devices feature advanced diagnostics and monitoring, including a power on self test (POST), which is performed by all SSA devices. Each port remains in a disabled state until it has successfully completed a beginning-communication process.

1.2 IBM 7190-100 SCSI Host to SSA Loop Attachment

The 7190 Model 100 is a SCSI to SSA converter. It provides the benefits of IBM′s SSA disk subsystems to systems that have a SCSI-2 fast wide differential controller. It is an external attachable device that does not require internal installation or configuration of host systems. It uses the SCSI bandwidth of 20 MB/s more efficiently by utilizing SSA devices that do not require arbitration. The 7190 Model 100 is host software independent and the SSA connectors are industry standard. It is supported on various Sun platforms. (Refer to Chapter 3, ″Planning and Design″ for exact system requirements.)

1.2.1 Description

Figure 2 shows a simple configuration using the 7190 Model 100.

The 7190 Model 100 is approximately 73 mm (2.9 in.) high, 162 mm (6.5 in.) wide, and 221 mm (8.8 in.) long. It is a self-contained unit with an internal power supply and a fan. The front panel consists of a power-on LED and a status LED. The rear panel has a Power switch, an AC power input mains, two industry standard SSA connectors, and a SCSI-2 fast wide differential connector.

1.2.2 System Configuration

Figure 2 represents the 7190 model 100 in a loop topology.

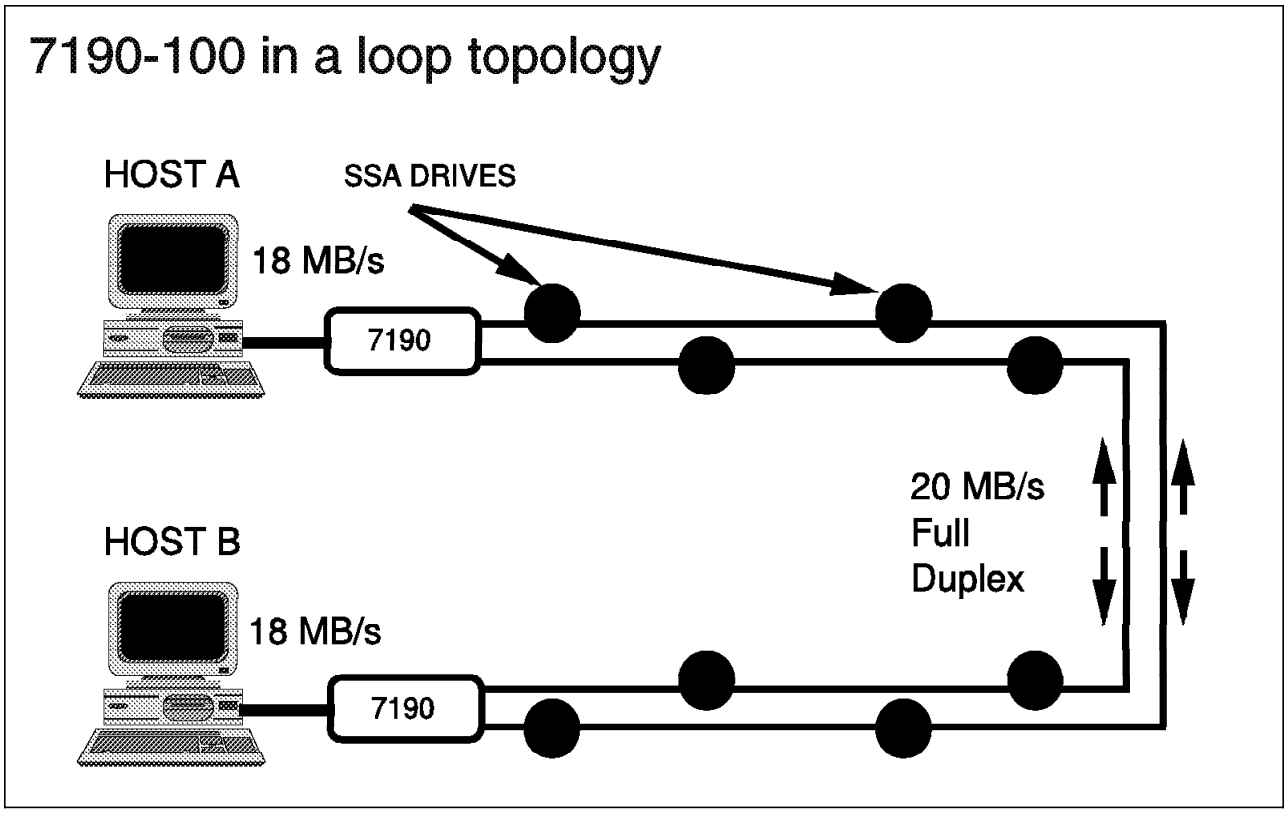

Figure 2. 7190-100 in a Loop Topology

The 7190 model 100 can attach to any SSA topology, but a loop is preferred. A maximum of four 7190 Model 100s can be attached to an SSA loop. Up to two 7190 model 100s can be attached to the same host system on the same loop or on separate loops. A loop containing 7190 Model 100s may not contain other types of initiators. A maximum of 48 drives can be attached to the loop. Optimal performance is obtained when each host system has a 7190 Model 100 that interfaces with the SSA loop. The 7190 model 100 maps attached SSA devices to SCSI ID/Logical Unit Number (LUN) pairs so that the host system can treat them as normal SCSI devices. It builds a table in nonvolatile memory that records information about SSA devices attached to the 7190 Model 100.

1.2.3 Features

The 7190 Model 100 has the following features:

- Sustained data rate of 18 MB/s
- Up to 1900 I/O operations per second
- Host software independence
- SCSI ID/LUN mapping
- Automatic loop configuration
- 25 m (82 ft) between devices (fiber optic extender can be attached between SSA DASD subsystems)
- Online service monitoring

1.3 IBM SSA Interface Controller for Sun SBus

The IBM SSA Interface Controller for Sun SBus is a single-card converter that enables the attachment of IBM′s 7133 and 7131-405 DASD subsystems to Sun servers running the Solaris operating system (refer to Chapter 3, ″Planning and Design,″ for exact system requirements). The SSA connectors are industry standard.

1.3.1 Description

Figure 3 shows the IBM SSA Interface Controller for Sun SBus.

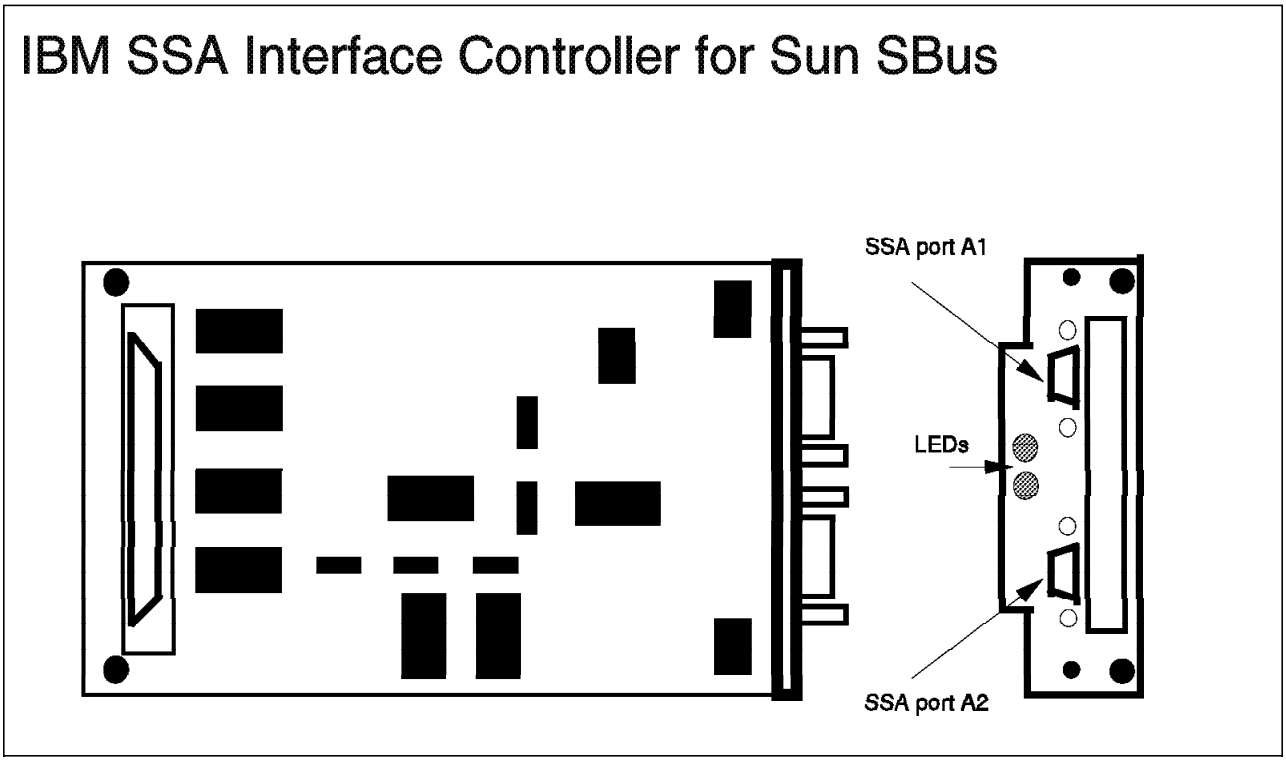

Figure 3. SBus Controller

The IBM SSA Interface Controller for Sun SBus is configured as a standard SBus adapter and is approximately 198 mm (7.9 in.) long and 84 mm (3.4 in.) wide. The controller consists of two SSA ports, a single-width SBus mezzanine card, and SBus SSA logic. The two SSA ports connect to industry standard SSA connectors on the bezel of the controller. There is also a warning LED and a status LED on the bezel.

1.3.2 System Configuration

Figure 4 on page 6 shows two IBM SSA Interface Controller for Sun SBus cards in the same system attached to two SSA loops.

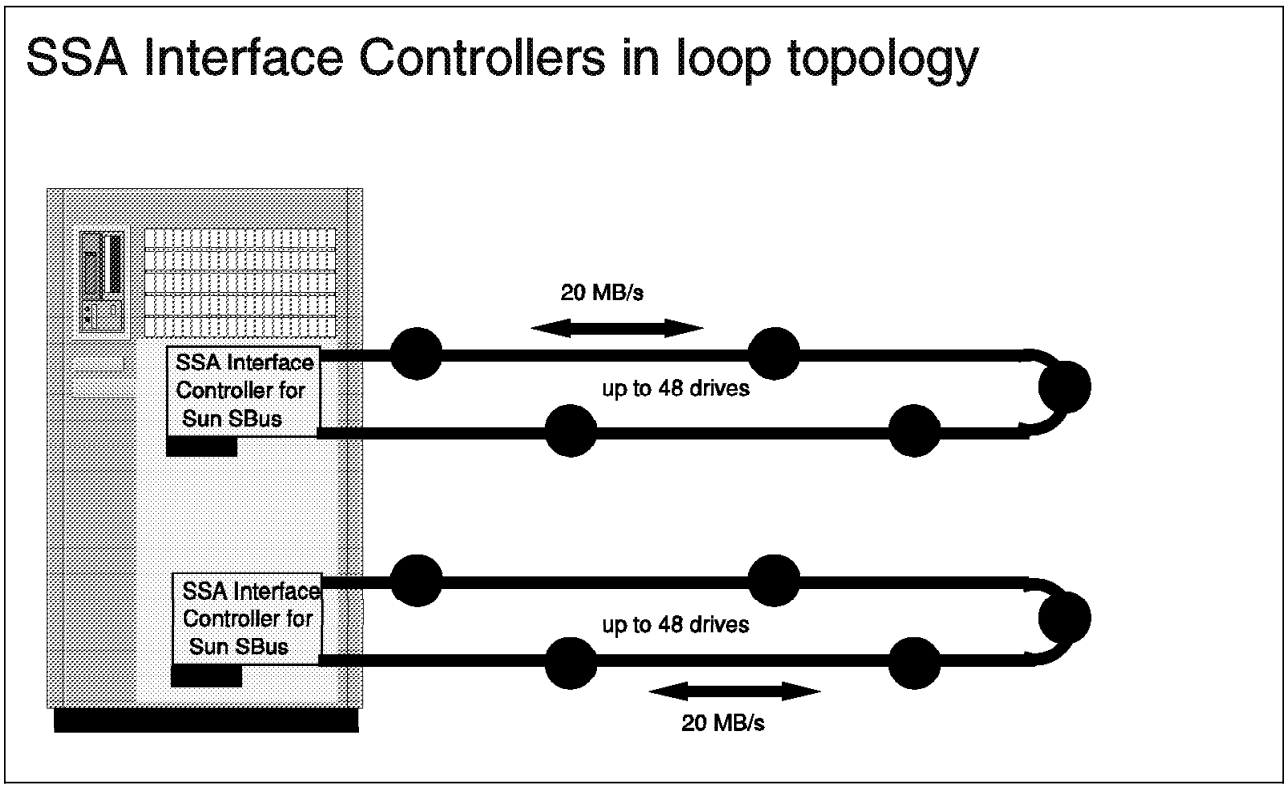

Figure 4. SBus Controllers in loop Topology

The IBM SSA Interface Controller for Sun SBus can be connected to any SSA topology, but the preferred configuration is a loop (figref refid=GAR03.). Up to two controllers per system are supported, provided that each controller is attached to a separate loop. One system can support up to 96 drives, as long as there are no more than 48 drives on each loop. The maximum configuration is four systems, eight adapters, and two loops. Each of the four systems can have up to two SBus adapters, provided the SSA Interface Controller cards in each system are attached to separate loops.

1.3.3 Features

The controller has the following features:

- Sustained data rate of 35 MB/s
- Up to 3000 4KB I/O operations per second
- Command queuing: 220 commands per SSA node
- 873 GB of serial storage
- 25 m (82 ft.) between devices (extendable with fibre optic extender)

1.4 IBM SSA DASD Subsystems

IBM has two SSA DASD subsystems, the IBM Multi Storage Tower Model 405 (7131-405) and the IBM 7133 Serial Storage Architecture Disk Subsystem (7133). They are attached to the host system by means of an SSA adapter, which for the purpose of this book, refers to either the 7190 model 100 or the IBM SSA Interface Controller for Sun SBus. The subsystems are independent of the SSA adapter and are easily portable. New technology can be added to the subsystem

by replacing the low-cost adapter with an updated version, which provides lower costs and investment protection.

1.4.1 IBM 7131 Multi storage Tower Model 405

The 7131 Model 405 is IBM′s entry level SSA disk subsystem. It is a low-cost cost solution that provides up to 45.5 GB of SSA storage. The tower uses serial loop architecture to ensure redundant paths to all disks.

Although the 7131-405 is called a *tower*, it is small enough to sit on a desktop. It consists of five hot-pluggable slots. A minimum of two slots are occupied with 2.2 GB, 4.5 GB, or 9.1 GB drives. The three additional slots can be filled at the time of purchase or as needed.

1.4.2 IBM 7133 Serial Storage Architecture Disk Subsystem

The 7133 provides up to 145 GB of storage. Its serial loop architecture ensures redundant paths to all disks. It is available in 48 cm (19 in.) rack format as the 7133 Model 010 and 7133 Model 020. The Model 500 and Model 600 are deskside units.

The 7133 Models 010 and 500 can be populated with 1.1 GB, 2.2 GB, and 4.5 GB drives. Models 020 and 600 can be populated with 2.2 GB, 4.5 GB, and 9.1 GB drives. Each individual drive is mounted in an auto-docking carrier for easy replacement.

1.5 Administrative Requirements

The administrator is responsible for the following activities:

- Installation of the SSA adapter (7190 Model 100 or IBM SSA Interface Controller for Sun SBus)
- Device configuration
- Controller software installation
- Maintenance There is a difference in maintenance practice between the 7190 Model 100 and the IBM SSA Interface Controller for Sun SBus. A service representative must order a replacement 7190 Model 100 and perform the installation, but customers can order a replacement SSA Interface Controller for Sun SBus card and perform the installation themselves.
- Adapter card replacement

Chapter 2. SBus Adapter and 7190 Technology

In this chapter we explore SSA and explain how it is implemented with the 7190 and the SBus adapter. We start with a look at the attributes of the architecture itself. Then we examine the SBus adapter and the 7190, two hardware solutions for using this powerful architecture in the Sun Solaris environment.

2.1 Serial Storage Architecture

SSA is an innovative, high-performance serial interface designed to connect disk drives, optical drives, tape drives, CD-ROMs, printers, scanners, and other peripherals to personal computers, workstations, servers, and storage subsystems. As a foundation for understanding the implementation of SSA on the SBus adapter and 7190, we'll explore the highlights of the interface and protocol.

2.1.1 Protocol Standards

The ANSI T10.1 draft standards define the SSA interface and protocol in a structure of three layers: a physical interconnect layer, a lower layer protocol (LLP), and an upper layer protocol (ULP). Figure 5 illustrates these layers.

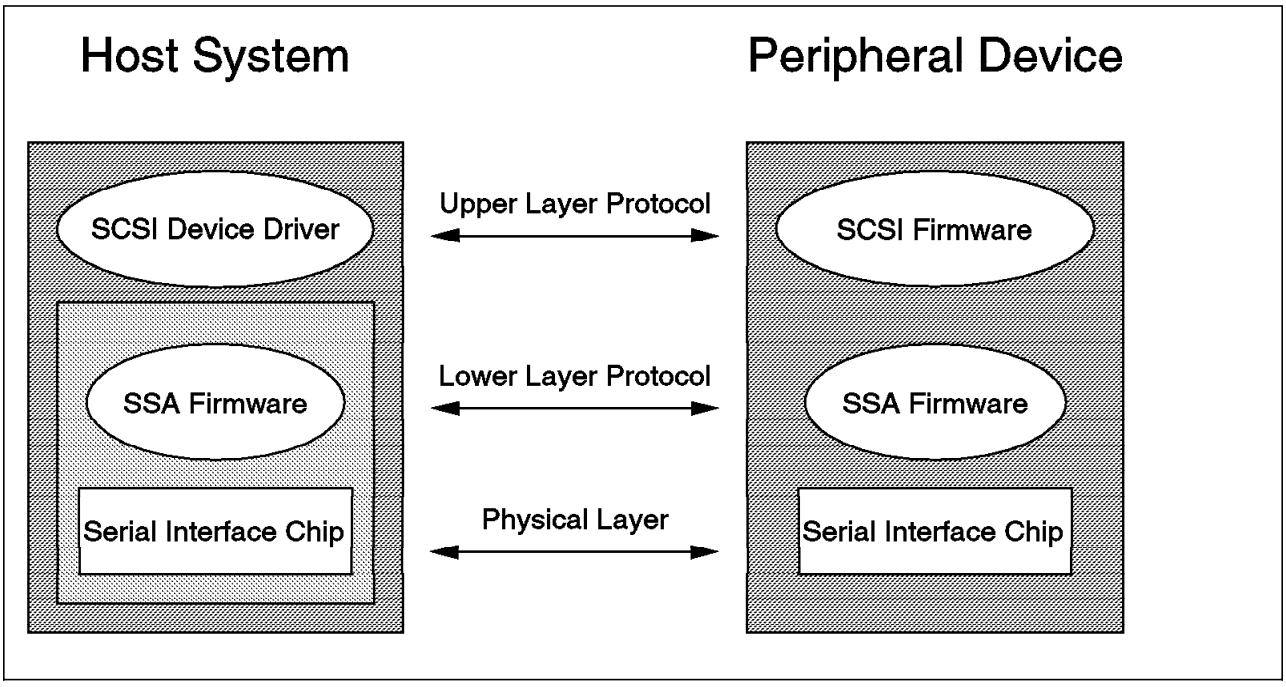

Figure 5. SSA Interface Layers.

2.1.1.1 Physical Interconnect Layer

The physical interconnect layer, also called the *physical layer*, is defined by the SSA-PH1 standard. The standard provides protocols to regulate such physical attributes as cables, connectors, transfer speeds, and data encoding. In essence, these are the attributes which describe how the parts are physically assembled and how they are used to transmit ones and zeros from one piece of hardware to another.

2.1.1.2 Lower Layer Protocol

The LLP, also known as the *transport layer*, is defined by the SSA-TL1 standard. This standard governs transport issues including link management, link protocol, framing, addressing, buffering, link resets, link configuration, and error recovery. These attributes enable the signaling capabilities provided by the physical interconnect layer to be used to reliably transmit groups of bytes (called *frames*) across a network (called a *web*) of SSA devices (called *nodes*) with low error rates and an error recovery mechanism. The frames are passed along the network from node to node, until they arrive at their destination. Each passage of a frame between two nodes is called a hop.

Transport layer frames vary in size from 7 to 139 bytes, and fall into three categories: control, privileged, and application. Control and privileged frames are used by the transport layer for link configuration and error recovery. Application frames contain either data or SSA Message Structures (SMSs) sent by the ULP. SMSs are used by the ULP for commands, data, status, and vendor-specific information. All frames use cyclic redundancy checks (CRCs) as an error-detection mechanism. Error recovery for corrupted frames is accomplished by retransmission, but with a raw error rate of less than 1 in 10^{13}, this is seldom necessary.

2.1.1.3 Upper Layer Protocol

The ULP is defined by the SSA-S2P standard. This standard details SCSI-2 as an upper layer protocol and describes how SCSI-2 is to be mapped to SSA using the transport layer. Mapping retains the SCSI functions associated with initiators, targets, and logical units, and enables SCSI commands, status, and messages to be enclosed in SSA application frames for transport across the SSA web. This SCSI mapping enables host computers and peripherals to communicate using SSA.

Note that the SSA physical and transport layers can support any ULP. SCSI was chosen becauuse of its popularity. Using SCSI as the ULP allows easy migration of existing products to SSA because much of the SCSI logic is retained, so code rewrites are minimized. SSA simply adds a layer of firmware for the transport layer.

2.1.1.4 Other Standards

There are other documents, published by the SSA Industry Association (SSA-IA), which define current implementations of SSA. SSA-IA/95PH is a combination of the ANSI documents SSA-PH1 and SSA-TL1. The SSA-IA document is nearly identical to its ANSI counterparts. SSA-IA/95SP is an almost identical copy of ANSI′s SSA-S2P. The SBus adapter and 7190 implement the protocols as defined in the SSA-IA documents.

The SSA protocol is being enhanced. Future versions of the ANSI standards will provide for modifications including increased link speeds, increased distances, and mapping of SCSI-3 as the ULP. The SSA-IA standards and ANSI standards are compatible with one another, and the current objective of the ANSI standards committee is to keep them compatible.

2.1.2 Spatial Reuse

An important attribute of SSA is the concept of spatial reuse. In a bus topology, data is broadcast across the entire bus. The entire bus is involved in only one transaction at a given moment in time. An example is Ethernet, where only one node transmits data on the network at a time while other nodes must wait until the network is quiet before transmitting.

In contrast, SSA implements the property of spatial reuse. Data is routed only over specifically designated links between the source and the destination. Thus multiple transfers can occur simultaneously at full bandwidth, provided the transfers use different links.

The difference between these two attributes can be envisioned by comparing a single-lane road to a freeway with two lanes in either direction. The one lane road can support one vehicle traveling in one direction at a given speed and time. In the same time frame, the freeway can support two vehicles traveling in either direction at the given speed.

2.1.3 Addressing Scheme for Simplicity

The addressing scheme of SSA greatly simplifies the task of physical configuration over that of SCSI. SCSI implements up to 16 devices on a bus, with addresses which are set by switches, jumpers, or software parameters. SSA defines an actual network topology. A network, also called a *web*, can be a string, a loop, or a collection of strings and loops interconnected by switches. Since this is a network topology disks, controllers, and other devices on the web are called *nodes*. A string can contain up to 129 nodes, a loop can contain up to 128 nodes, and a complex web connected by switches can contain nearly 2 million nodes.

With this potential for large numbers of nodes on an SSA web, identification could be confusing. Because of this, the architecture uses unique identifiers (UIDs) to uniquely identify every possible node in the SSA web. UIDs consist of an 8-byte field: 2 reserved bytes (containing zeros), a 3-byte organizationally unique identifier (OUI), and a 3-byte product identifier assigned by the assignee of the OUI. The OUI is assigned by the Institute of Electrical and Electronic Engineers (IEEE). Figure 6 on page 12 shows an example of an SSA UID.

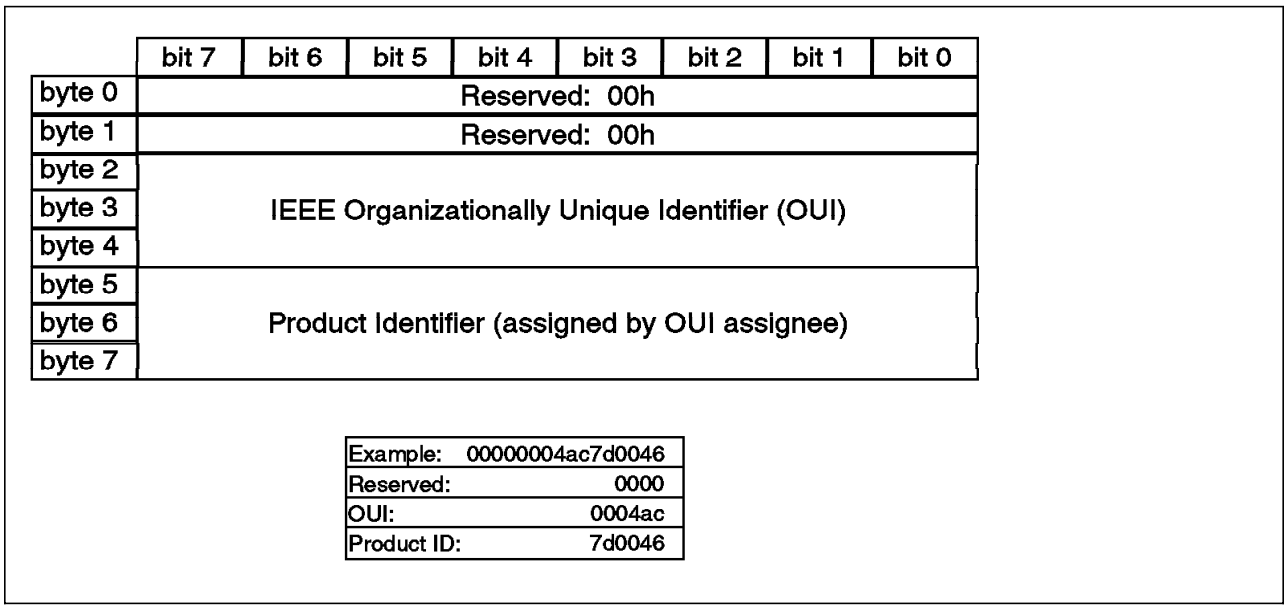

Figure 6. SSA Unique Identifier.

One benefit of the UID is that it guarantees that each node on the web is uniquely identified. Another benefit is that the UID allows reconfiguration of the web to be unobtrustive: devices can be added to or removed from an SSA web without powering down the other components on the web. This is permitted by the guarantee of address uniqueness and the ability of the LLP to handle changes in an active web. Removing a device from an SSA loop simply divides the loop into two strings, which continue to operate without the need for any special bypass hardware. Compare this with SCSI, which requires that the bus be powered off to add or remove devices or to change an address of a particular device. SCSI also has the potential for misconfiguration of the bus by accidentally placing two devices with the same address on the same bus.

2.1.4 Fairness Algorithm

To ensure that no one device dominates the loop, SSA implements a fairness algorithm. Simply stated, the algorithm involves the circulation of tokens (called *SAT* tokens, for satisfaction). The tokens circulate in both directions around a loop, or from end to end on a string.

The principle of the fairness algorithm is that each node will, over time, accumulate a queue of I/O, either data or commands and responses, to transmit. Typically, if a node is not passing messages from the adjacent node on one side to the adjacent node on the other side, it sends its I/O on the appropriate link, which would be quiet. If the node is busy passing traffic from one side to the other and still wants to transmit I/O, it queues the I/O. In theory, this poor node could be stuck passing traffic for other nodes and unable to transmit its own.

Enter the SAT token. One token circulates in each direction around the loop (or along the string). Simply stated, if the node receives a SAT token and has a queue of I/O waiting to transmit, it is allowed to transmit the I/O at that time. This transmission is allowed up to a level at which the node is considered “satisfied.” At this point, the node must again wait until either the link is quiet or it receives another SAT token.

2.1.5 Terminology

There are a few terms within SSA that you should be familiar with:

- A node is a system, controller, or device on an SSA web with one or more serial links.
- A configutor (yes, that is how it is spelled) is a node that knows the topology of its associated SSA web. It is often referred to as an *initiator*, but configutor is the proper term for a node at the LLP (transport layer) of SSA.
- A responder is a node that does not know the topology of the its associated SSA web. It is often referred to as a *target*, but responder is the proper term for a node at the transport layer of SSA.
- An initiator is a component within a device that determines what task needs to be executed and which target will perform the task. An initiator will originate operations, create frames of data to send across an SSA link, and will confirm that the target performed its assigned task. The term initiator generally refers to devices at the ULP (SCSI) layer of SSA.
- A target is a component within a device that performs operations and returns SCSI status to the initiator. A target is a destination for frames and commands, and may have up to 128 logical units (LUNs) attached to it. The term *target* generally refers to devices at the ULP (SCSI) layer of SSA.
- A logical unit is a physical or virtual peripheral device which is addressable through a target. It can be a portion of a device (often a disk). Each device will have one or many LUNs associated with it. LUN is another ULP term.

2.2 IBM SSA Interface Controller for Sun SBus Technology

The IBM SSA Interface Controller for Sun SBus is an SBus adapter card that attaches a single SSA loop to a Sun SBus. The SBus adapter brings the speed and dependability benefits of IBM′s SSA storage to selected Sun platforms in the form of a native SSA interface. Figure 7 on page 14 shows a simple SSA implementation using the SBus adapter. A summary of SBus adapter features is shown in Figure 8 on page 14. Operating specifications for the SBus adapter are shown in Figure 9 on page 15.

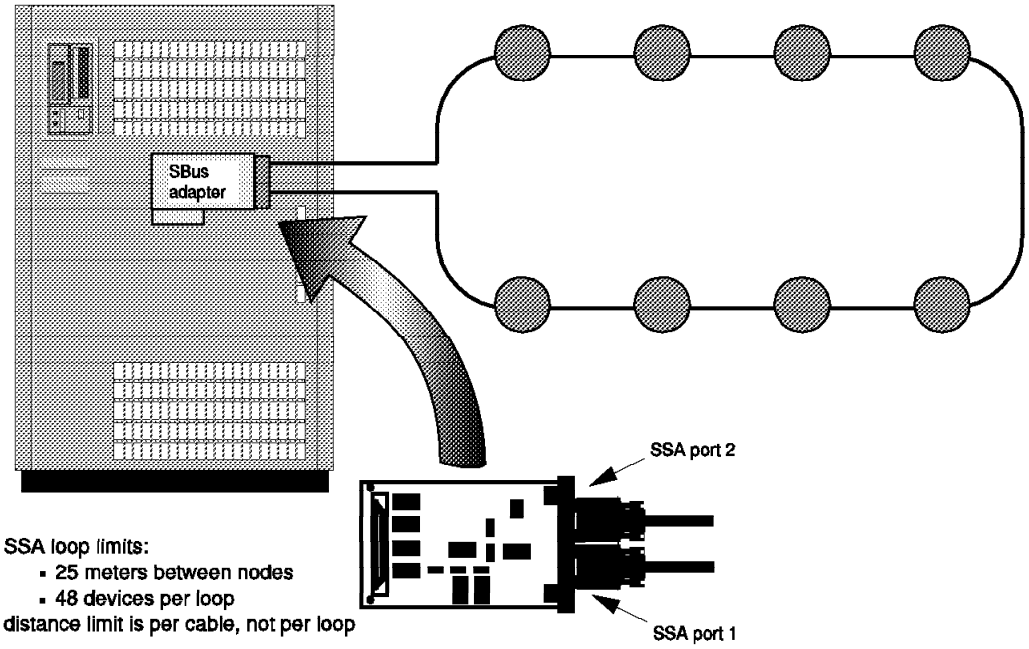

Figure 7. Simple Loop Using SBus Adapter.

Summary of SBus Adapter Features

SSA bandwidth:	80 MB/s
Sustained transfer rate:	up to 35 MB/s
Sustained transaction rate:	up to 3000 I/O operations/s
Adapters per host:	two
Adapters per SSA loop:	four (see Note 1)
Adapters per host in same loop:	two
Drives per SSA loop:	48

Note 1: All adapters in a single loop must be SBus Adapters

Figure 8. Summary of Features: SBus Adapter.

Power consumption:	8.5 W (estimated)
Power connection:	obtained from SBus interface
Airflow requirement:	0.5 m3/s (minimum)
Form factor:	single-width SBus card
Dimensions:	146.7 mm x 83.8 mm (5.776 x 3.300 in)
Thickness:	1.6 mm (.062 in)
Weight	116 g (4.09 oz)
Operating environment:	
Air temperature	10 to 40 degrees C (50 to 104 degrees F)
Relative humidity	8 to 81% noncondensing
Maximum wet bulb temperature	27 degrees C (80 degrees F)
Altitude	0 to 2133 m (0 to 7000 ft)

Figure 9. Operating Specifications: SBus Adapter.

The SBus adapter incorporates several components. In this section, we examine the functions of the major components.

2.2.1 80960JF Microprocessor

The Intel i960 processor is a 32-bit RISC microprocessor capable of executing one instruction per clock cycle. The 80960JF-25 implementation runs at 25 MHz, providing a performance of approximately 23 million instructions per second (MIPS). The processor executes the adapter firmware, responds to interrupts from the serial interface chip, and coordinates transactions between the serial interface chip and the SBus interface. The 2 KB write-through data cache and 4 KB instruction cache increase performance by reducing memory accesses required for instruction fetches and data reads.

2.2.2 Serial Interface Chip

The IBM serial interface chip (SIC) is the nerve center of the SSA connectivity of the adapter. The chip is compliant with SSA Industry Association (SSA-IA) standards SSA-IA/95PH and SSA-IA/95SP.

Implemented in 3.3 V complementary metal oxide semiconductor (CMOS) hardware, the SIC is a single-chip solution that integrates the SSA protocol and two SSA ports in a unit. The unit can be used for either configutor or responder devices, as well as switches. The SIC interfaces with the attached SSA network using two 20 MB/s serial ports running in full duplex mode. This affords a total bandwidth of 80 MB/s, bidirectional, full duplex. Future increases in this bandwidth are easily achievable through changes in clock speed.

By comparison, the closest SCSI equivalent to the SIC′s functionality requires several components, and increases in bandwidth are not so simple. Parallel bus issues such as clock skew and line-length limitations make it difficult to accelerate data transfer through increased clock speed. Increasing data transfer

by widening the SCSI bus means increasing the number of wires and contacts in the SCSI cable beyond the current 50 or 68, making the cabling more expensive and more difficult to work with than the four-wire external or six-wire internal SSA cables.

2.2.3 Flash Memory

The 512 KB flash memory module stores adapter functional code, firmware, and the power-on self-test (POST) code. The firmware contained in flash memory can be updated through a host download procedure.

2.2.4 Fast Static Random Access Memory

The logic card includes 512 KB of fast static random access memory (SRAM), which acts as a buffer for the SIC. It also provides an execution area for the adapter firmware, once the firmware is copied from flash memory.

2.2.5 SSA Connectors

The SBus adapter implements the physical connection to a single SSA loop or two strings with two nine-pin micro-D connectors. SSA provides for automatic bus termination, so users do not have to add or remove termination devices from the SSA connections.

2.2.6 SBus Connector

Located on the underside of the SBus adapter is a 96-pin, high density connector. This connector fits into a host system SBus slot that is compatible with IEEE 1496-1993.

2.2.7 Light Emitting Diodes

Two light-emitting diodes (LEDs), one green and one yellow, are positioned between the SSA ports. The LEDs provide an external indication of adapter status. These status indications are summarized in Figure 10 on page 17.

LED Activity		Explanation	Action
Yellow ON	Green flickers	Adapter initialization.	None
Yellow ON	Green OFF or ON	A fatal error has occurred during Adapter initialization.	Use SSAU to run adapter diagnostics. If error persists replace adapter.
Yellow ON	Green flashes in patterned sequence	A fatal error has occurred during Adapter operation; the green LED flashes in a four-digit error code.	Record LED sequence. Use SSAU to run adapter diagnostics. If error persists, replace adapter.
Yellow ON	Green flashes once every few seconds	Adapter is running; both ports are disconnected.	Check cabling. Check nodes connected to adapter.
Yellow OFF	Green flashes steadily	Adapter is running; one port is connected.	Only one port is active. Check cabling. Use SSAU web map function to identify inactive port.
Yellow OFF	Green ON	Adapter is running; both ports are connected.	None
Yellow OFF	Green OFF	No power to the Adapter	Turn ON the system, or check power supply.

Figure 10. Summary of LED Activity: SBus Adapter.

2.3 IBM 7190 Technology

The IBM 7190 Model 100 is an interface adapter that attaches a single SSA loop to a fast wide differential SCSI bus. The 7190 brings the speed and dependability benefits of IBM′s SSA storage to host platforms that do not have a native SSA interface. Because the host system′s native SCSI is used, the 7190 offers the advantage of not having to modify host system software or install additional device drivers. Figure 11 on page 18 shows a simple SSA implementation using the 7190. A summary of 7190 features is shown in Figure 12 on page 18. Operating specifications for the 7190 are shown in Figure 13 on page 19.

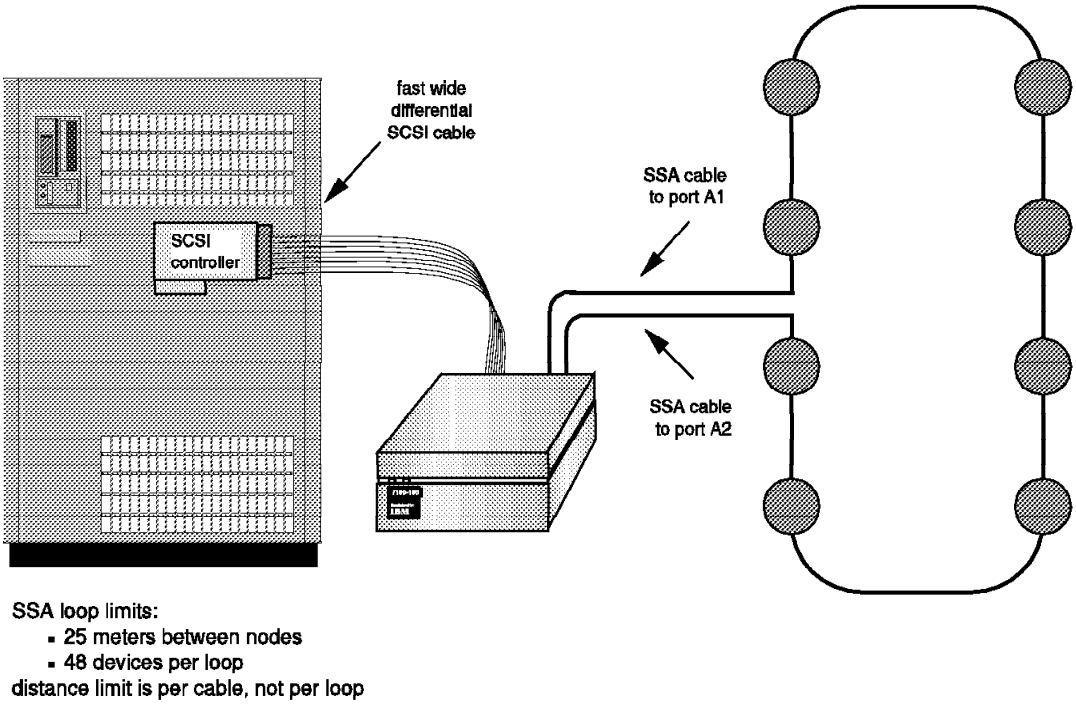

Figure 11. Simple Loop Using 7190.

Summary of 7190 Features

SSA bandwidth:	80 MB/s
Sustained transfer rate:	up to 18 MB/s (SCSI limit)
Sustained transaction rate:	up to 1900 I/O operations/s (SCSI limit)
Adapters per host:	limited by number of SCSI adapters
Adapters per SSA loop:	four (see note 1)
Adapters per host in same loop:	two
Drives per SSA loop:	48

Note 1: All adapters in a single loop must be 7190s

Figure 12. Summary of Features: 7190.

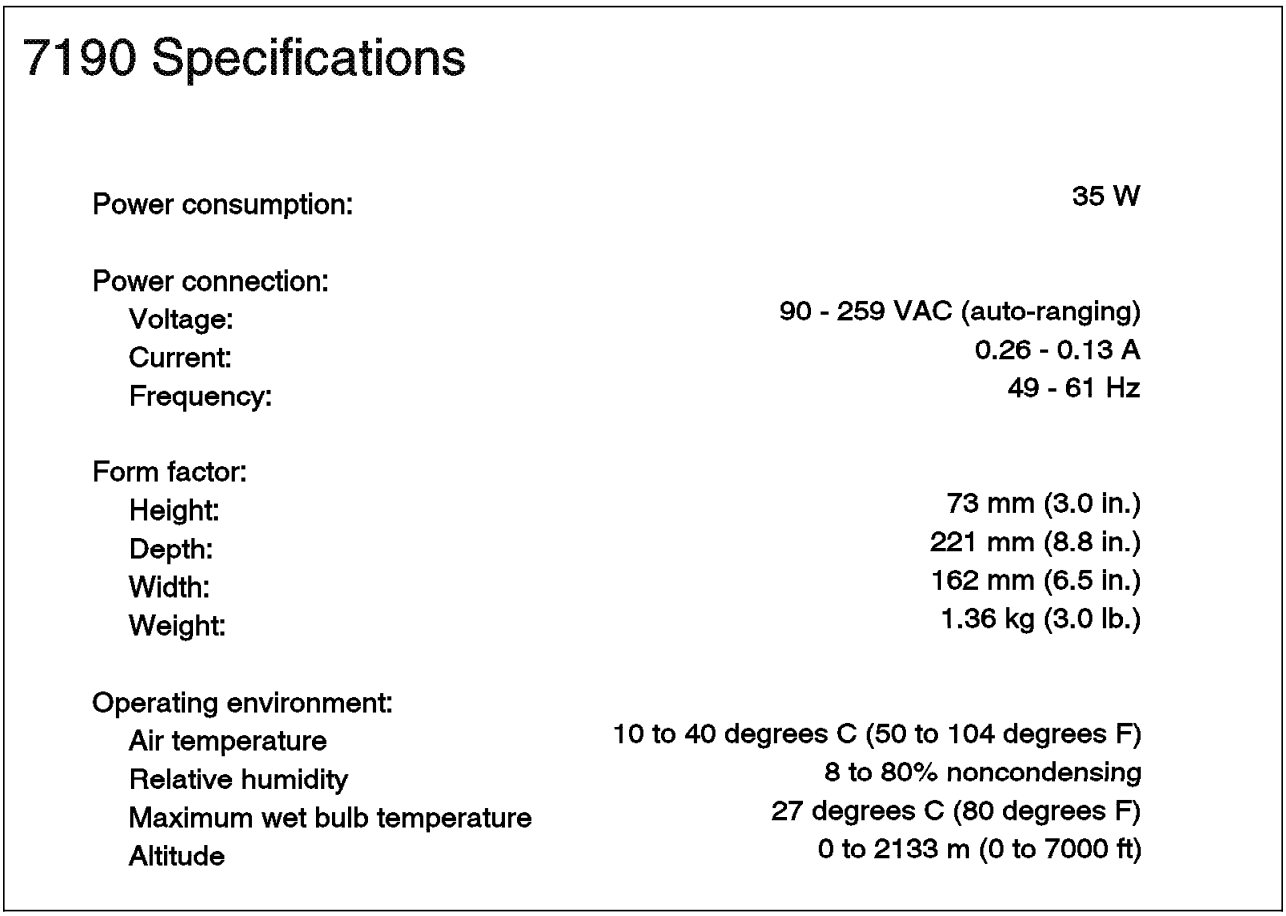

Figure 13. Operating Specifications for the 7190.

The 7190 incorporates a number of different physical components, which we now examine in some detail. Figure 14 on page 20 shows the external components of the 7190.

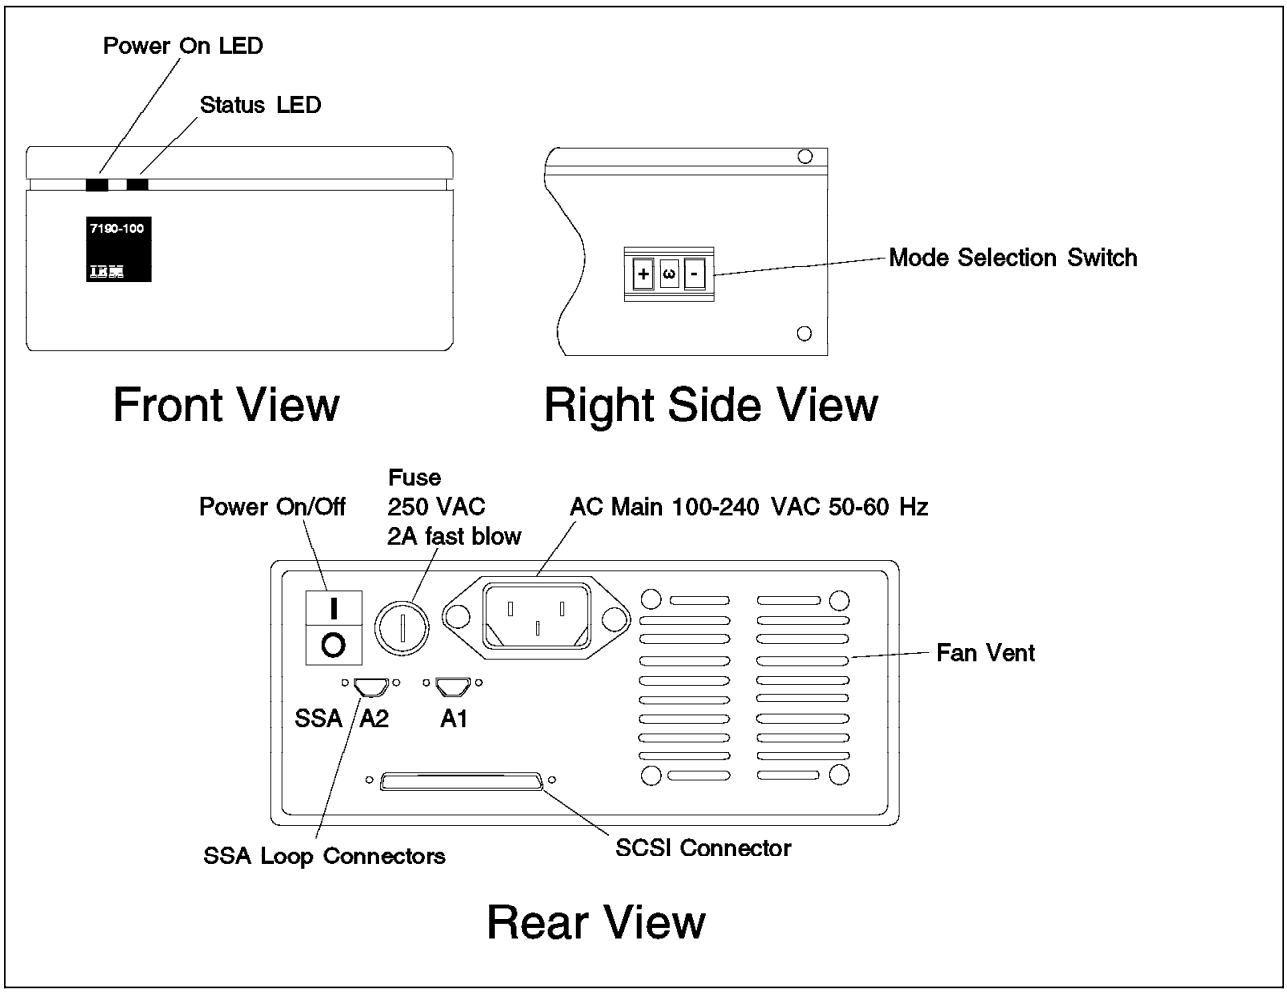

Figure 14. IBM 7190 External Components.

2.3.1 80960Dx Microprocessor

The Intel i960Dx processor is a 32-bit RISC microprocessor capable of executing one instruction per clock cycle. The processor executes the 7190 firmware, responds to interrupts from the SIC, and coordinates transactions between the SSA serial interface chip, the buffer memory, and the SCSI controller. The 2 KB write-through data cache and 4 KB instruction cache increase performance by reducing memory accesses required for instruction fetches and data reads.

2.3.2 Symbios Logic 53C825A

The Symbios Logic 53C825A SCSI controller provides connectivity that complies with the ANSI X3.131-1994 specification for SCSI-2. The controller provides a maximum burst throughput of 20 MB/s on a fast wide differential interface.

2.3.3 IBM Serial Interface Chip

The IBM SIC is the nerve center of the SSA connectivity of the adapter. The chip complies with SSA-IA standards SSA-IA/95PH and SSA-IA/95SP.

Implemented in 3.3 V CMOS hardware, the SIC is a single-chip solution that integrates the SSA protocol and two SSA ports in a unit. The unit can be used for either initiator or target devices (in SSA, they′re called *configutor* and

responder, respectively), as well as switches. The SIC interfaces with the attached SSA network using two 20 MB/s serial ports running in full duplex mode. This affords a total bandwidth of 80 MB/s, bidirectional, full duplex. Future increases in this bandwidth are easily achievable through changes in clock speed.

By comparison, the closest SCSI equivalent to the SIC′s functionality requires several components, and increases in bandwidth are not so simple. Parallel bus issues such as clock skew and line-length limitations make it difficult to accelerate data transfer through increased clock speed. Increasing data transfer by widening the SCSI bus means increasing the number of wires and contacts in the SCSI cable beyond the current 50 or 68, making the cabling more expensive and more difficult to work with than the four-wire external or six-wire internal SSA cables.

2.3.4 SCSI Connector

The 7190 connects to the SCSI bus, using a single 68-pin high-density mini-D connector. The use of a single connector ensures that the 7190 is the last device on the SCSI bus. This is because the 7190 provides the required differential signal termination internally. The user must ensure that correct differential termination is supplied on the host SCSI adapter end.

2.3.5 SSA Connectors

The 7190 adapter implements the physical connection to a single SSA loop or two strings with two 9-pin micro-D connectors. SSA provides automatic bus termination, so users do not have to add or remove termination devices from the SSA connections.

2.3.6 SSA to SCSI ID/LUN Mapping

The 7190 maps the attached SSA devices to SCSI target ID and LUN pairs. The maps are stored in a table in nonvolatile memory within the 7190. In the table, each disk is identified by its SSA UID, which is indexed to the assigned SCSI ID/LUN pair.

When the map is first established, the order of selection of SSA disks starts with the device nearest port 1 of the 7190 and continues to the device farthest from port 1. If the topology is a loop, the device farthest from port 1 is the device nearest port 2. If the topology is a string, selection continues with the device nearest port 2 to the device furthest from port 2. An example of this ID/LUN mapping is shown in Figure 15 on page 22.

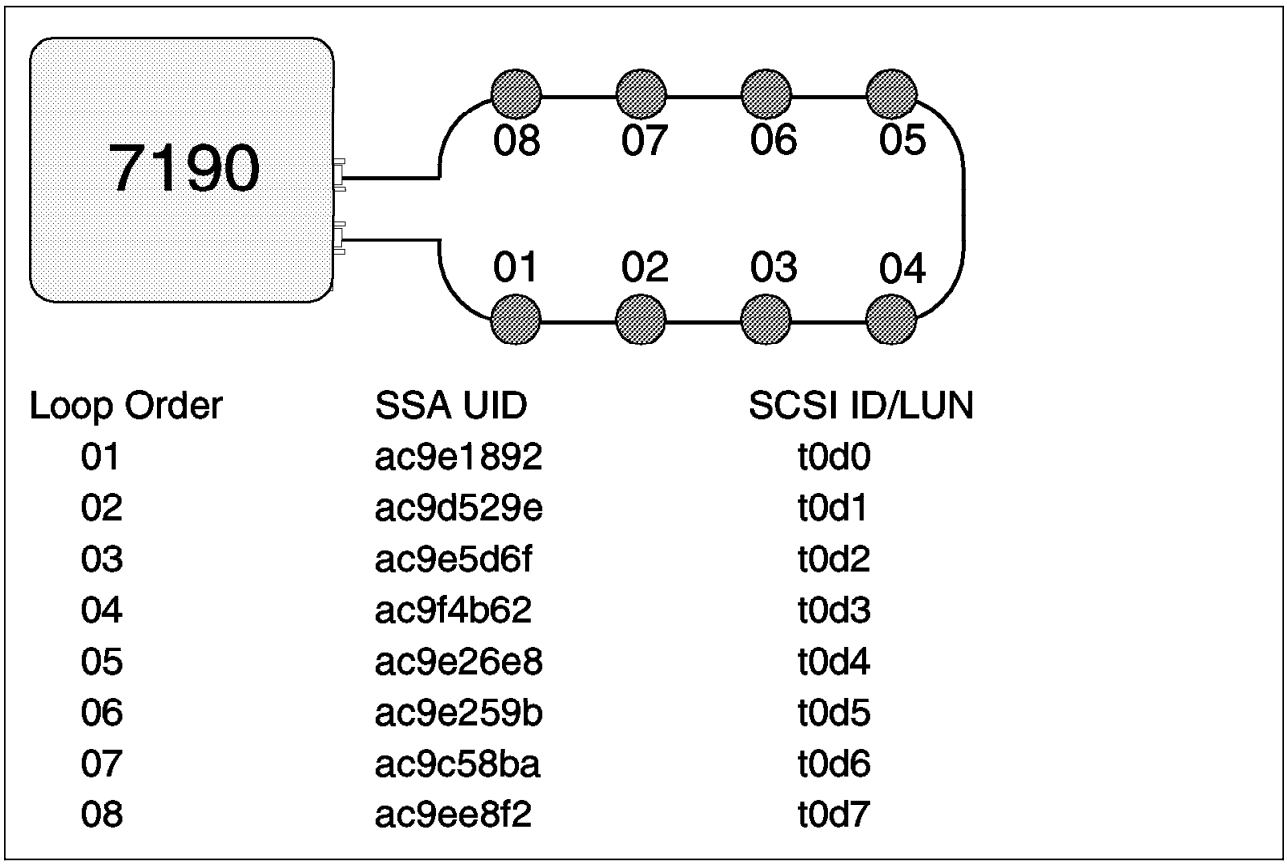

Figure 15. Mapping SSA UIDs to SCSI Target ID/LUN Pairs.

Mapping is performed in two different styles: ID priority and LUN priority.

2.3.6.1 Mode 0: ID Priority Mapping

The first style of SCSI ID/LUN mapping is referred to as ID priority mapping, fast wide SCSI, with device spin up. This mode of operation is achieved by selecting 0 (zero) on the 7190 mode select switch.

ID priority mapping assigns the disks to SCSI IDs of 0 through 15, except ID 7. The numbering sequence starts with ID 0 LUN 0 and increments the ID number by 1 through ID 6 LUN 0. Numbering then skips to ID 8 LUN 0, ID 9 LUN 0, continuing to ID 15 LUN 0.

The sequence repeats, starting with ID 1 LUN 1, ID 2 LUN 1, and so forth. This method provides a maximum of 120 SCSI ID/LUN pairs, leaving ID 7 for the SCSI adapter. Figure 16 on page 23 illustrates the concept of ID priority mapping.

LUN 0	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 1	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 2	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 3	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 4	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 5	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 6	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15
LUN 7	ID 0	ID 1	ID 2	ID 3	ID 4	ID 5	ID 6	(reserved)
	ID 8	ID 9	ID 10	ID 11	ID 12	ID 13	ID 14	ID 15

In ID priority mapping, all 15 IDs are assigned a LUN before proceeding to the next LUN number. The first combination assigned is ID 0 LUN 0, then ID 1 LUN 0, moving across the row to ID 6 LUN 0. Following ID 6 LUN 0, move to the next row starting with ID 8 LUN 0, moving across row to ID 15 LUN 0. Following ID 15 LUN 0, move to the next row starting with ID 0 LUN 1, moving across the row to ID 6 LUN 1. Move to the next row for IDs 8 through 15 LUN 1, continue with rows for LUNs 2 through 7.

Figure 16. ID Priority Mapping.

2.3.6.2 Modes 2 and 3: LUN Priority Mapping

The second style of SCSI ID/LUN mapping is known as LUN priority. This mode of operation is achieved by selecting Mode 2 (LUN priority without device spin up) or Mode 3 (LUN priority with device spin up) on the 7190 mode select switch.

LUN priority mapping uses SCSI IDs 0 through 15, except for IDs 6 and 7. The numbering sequence starts with ID 0 LUN 0 and increments the LUN number by 1 through ID 0 LUN 7. The sequence repeats, starting with ID 1 LUN 0, ID 1 LUN 1, and so forth, continuing to ID 6 LUN 7. The sequence then skips to ID 8 LUN 0 and continues to ID 15 LUN 7. This method provides a maximum of 112 SCSI ID/LUN pairs, and leaves IDs 6 and 7 for the SCSI adapter and a possible spare. Two adapters could be placed on the bus, although this is not recommended for performance reasons. Figure 17 on page 24 illustrates the concept of LUN priority mapping.

ID 0	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 1	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 2	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 3	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 4	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 5	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 6	reserved for SCSI adapter							
ID 7								
ID 8	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 9	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 10	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 11	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 12	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 13	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 14	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7
ID 15	LUN 0	LUN 1	LUN 2	LUN 3	LUN 4	LUN 5	LUN 6	LUN 7

In LUN priority mapping, all eight LUN numbers are assigned to a given ID before incrementing the ID.
The first combination assigned is ID 0 LUN 0, moving across the row to ID 0 LUN 7.
Following ID 0 LUN 7, move to the next row starting with ID 1 LUN 0, moving across row to ID 1 LUN 7.
Following ID 1 LUN 7, move to the next row and repeat for IDs 2 through 5, then IDs 8 through 15.

Figure 17. LUN Priority Mapping.

Mode 3, LUN priority with device spin up, is the factory default setting, and is the setting most commonly used. Generally, ID priority is used only when more than 112 ID/LUN pairs are required.

2.3.7 Unarbitrated Path for Bandwidth

The 7190 provides an unarbitrated path to the SSA devices on the attached loop by appearing to the host adapter as a single SCSI target. This important design feature provides a much higher utilization of bandwidth than a comparable SCSI disk subsystem, because the only arbitration on the bus is between the SCSI controller and the 7190. A SCSI disk subsystem requires arbitration between all disks on the bus and the controller, which costs time. Moreover, the 7190 supports a wide range of advanced SCSI features to enhance performance, including tagged command queuing, scatter/gather, disconnect/reconnect, and both asynchronous and synchronous data transfers. These features allow such things as optimized ordering of commands and tuning of the SCSI bus. The net effect is that the 7190 is easily used as a direct replacement for SCSI devices, offering the performance benefits of SSA.

Coupled with high-performance disk drives, the 7190 delivers an I/O rate of up to 1,900 operations/second, with a data throughput rate of up to 18 MB/s.

2.3.8 Data Integrity

We believe that data integrity is, by far, the most crucial aspect of any storage subsystem. The 7190 was designed not only for performance and flexibility of attachment but also to ensure the highest data integrity possible. On the SCSI bus, parity is used to protect data traveling between the SCSI host adapter and the 7190. If a parity error is detected by the receiving device, the receiving device will request the sending device to retransmit the message. The SSA protocols, implemented on the 7190, protect data on the SSA loop by use of a

cyclic redundancy check (CRC). The CRC is implemented at the level of 4 bytes per SSA frame, with the frame not exceeding 139 bytes (including CRC). If a CRC error is detected by the receiving node, the receiving node will not acknowledge the affected frame and the sending node will retransmit it.

Chapter 3. IBM 7190 and SBus Adapter Planning and Design

In this chapter,we discuss elements that are important to designing your storage solution to take maximum advantage of the performance and scalability of SSA using either the SBus adapter or the IBM 7190. We introduce the host platform hardware and software prerequisites and explain how to arrange your storage system to maximize availability and performance. We also show you how to plan growth into your design and cover other considerations unique to SSA.

3.1 Host Platform Hardware Prerequisites

In this section, we discuss the hardware required for a supported 7190 or SBus adapter configuration in a Sun system environment.

3.1.1 Supported Models: SBus Adapter

The SBus adapter is designed for direct, high-speed connection of SSA storage systems to Sun SPARC server platforms. This native attachment is supported on the following systems:

Sun SPARCserver 20

Sun CPUs with 25 MHz SBuses are required for SPARCserver 20 systems. Other SPARCserver 20 models and clones have different bus timing which creates interface problems.

Sun SPARCserver 1000

Sun SPARCserver 2000

Sun SPARCserver 1000E

Sun SPARCcenter 2000E

Sun Ultra Enterprise 2, 3000, 4000, 5000, and 6000

The SBus adapter requires one 32-bit SBus-type adapter card slot compatible with the IEEE 1496-1993 Standard. Using the adapter requires OPEN BOOT BIOS firmware 2.x or above on any Sun platform. FORTH commands available on the SBus adapter were written for these versions. Earlier versions of OPEN BOOT BIOS have SBus firmware levels which are not compatible with the latest SBus cards, including the SSA SBus adapter.

3.1.2 Supported Models: 7190

The 7190 is designed for diversity of attachment across multiple hardware platforms. You can attach the 7190 to selected systems by using a fast wide differential adapter that is SCSI-2 compliant according to ANSI specification X3.131-1994. IBM supports the 7190 on a broad cross-section of Sun host platforms, including the following systems:

SPARCstation 10 and 20

SPARCserver 1000 and 1000E

SPARCcenter 2000 and 2000E

Ultra Enterprise 2, 150, 3000, 4000, 5000, and 6000

Attachment of the 7190 to a Sun platform must be through the Sun Microsystems, Inc. fast wide differential SCSI adapter, model X1062A.

3.1.3 Non-SSA System Disk Required for SBus Adapter and 7190

Neither the SBus adapter nor the 7190 support the use of SSA subsystems to boot the Solaris operating system. You must provide a non-SSA boot device that is supported for the specific hardware platform and Solaris version you are using. By keeping the SSA subsystem separate from the boot device, fault isolation between the SPARC system configuration and the SSA subsystem is greatly simplified. The 7190 service functions application and the SBus adapter service aid utility (SSAU) are normally located on the system disk and can be run without the SSA subsystem. In addition, service contract support is easily distinguished between the SPARC system and the SSA subsystem.

3.2 Host Platform Software Prerequisites

In this section, we discuss the software requirements for integrating a 7190 or SBus adapter into a Sun configuration.

3.2.1 Software Required for SBus Adapter

To use the SBus adapter supported systems, you need Version 2.4 or 2.5.1 of the Solaris operating system. These were the Solaris versions available when the SBus adapter became publicly available.

Each Solaris system must be configured with the latest patches recommended by Sun Microsystems. The minimum patch levels required for use of the SBus adapter are shown in Figure 18 on page 29 for Solaris Version 2.4, and Figure 19 on page 29 for Solaris Version 2.5.1. The patches listed are provided by Sun to protect customers from known Solaris problems. The SBus adapter was rigorously tested using these minimum levels, and the reliability of the adapter on systems below these levels is unpredictable.

3.2.2 Software Required for 7190

To use the 7190 on SuperSPARC systems (SPARCstation, SPARCserver, and SPARCcenter) you need Version 2.4 of the Solaris operating system with the patches listed in Figure 18 on page 29, Version 2.5 of Solaris, or Version 2.5.1 of Solaris with the patches listed in Figure 19 on page 29. To use the 7190 on UltraSPARC systems (Ultra Enterprise) you need Solaris Version 2.5.1 with the patches listed in Figure 19 on page 29. The patches listed are provided by Sun to protect customers from known Solaris problems. The 7190 was rigorously tested using these minimum levels, and the reliability of the 7190 on systems below these levels is unpredictable.

3.2.3 Other Requirements

Use of more than 32 drives attached to an SBus adapter or 7190 on all Solaris 2.4 systems requires patch 102283. This patch overcomes the SunOS 5.4 (Solaris 2.4) disks program limitation of 32 drives per controller.

Higher revisions of some or all of the listed patches may be available. In all cases, each Solaris system must be configured with the latest patches recommended by Sun Microsystems to provide protection from known Solaris problems. For the latest patch information or to download patches, see the

SunSolve Web pages listed << hdref reference to paragraph 4.1.1 ″Preparing for Installation″ >>.

101753-23	101880-10	101907-14	101922-16	102001-13	102037-01	102112-01
101829-01	101902-09	101920-03	101923-09	102011-03	102044-01	102283-01
101878-14	101905-01	101921-20	101945-46	102035-01	102062-03	102509-07
101879-02						

Figure 18. Solaris 2.4 Minimum Patch Levels. These are the minimum patch levels required to use the SBus adapter or 7190 with Solaris Version 2.4.

103558-07	103686-02	103847-02	104081-01	104283-01	104484-01	104635-01
103566-11	103690-03	103857-04	104093-02	104287-01	104489-01	104637-01
103582-08	103693-01	103866-02	104094-01	104297-01	104490-01	104642-01
103591-06	103696-02	103867-01	104162-02	104312-03	104516-01	104644-01
103594-09	103699-01	103879-04	104166-01	104317-01	104533-02	104650-01
103597-01	103734-01	103891-02	104182-01	104331-01	104538-01	104652-01
103600-12	103738-03	103900-01	104200-01	104334-01	104560-01	104654-01
103603-04	103743-01	103901-03	104201-01	104338-01	104576-01	104656-01
103606-02	103750-02	103903-02	104208-01	104361-01	104578-01	104658-01
103612-23	103770-02	103911-01	104212-03	104380-03	104580-02	104664-01
103618-02	103785-01	103928-01	104220-01	104389-01	104594-01	104672-01
103621-01	103793-01	103934-04	104227-01	104406-01	104595-02	104692-01
103622-05	103794-03	103948-01	104236-01	104433-02	104596-03	104696-01
103623-02	103795-02	103955-02	104239-01	104440-01	104604-02	104706-01
103627-01	103796-08	103959-03	104245-03	104443-02	104605-01	104708-01
103630-05	103801-04	103981-04	104246-03	104447-01	104607-01	104710-01
103640-08	103810-03	103995-01	104259-02	104453-01	104608-01	104736-01
103643-04	103817-01	104010-01	104263-01	104463-01	104609-01	
103663-07	103821-01	104029-01	104266-01	104475-01	104613-01	
103669-03	103834-01	104066-02	104277-01	104478-01	104615-01	
103680-01	103839-02	104072-01	104280-01	104482-01	104628-01	

Figure 19. Solaris 2.5.1 Minimum Patch Levels. These are the minimum patch levels required to use the SBus adapter or 7190 with Solaris Version 2.5.1.

3.3 Connectivity and Cabling

Connecting SSA storage to Sun systems through the 7190 or the SBus adapter can be very simple or somewhat complex, depending on your particular application needs and the size of your installation. Before tackling complex installations, let′s take a look at the basic building blocks of connectivity.

3.3.1 Loops Recommended

In even the most simple storage application, we recommend connecting the storage subsystem to the host system in a loop topology. Figure 20 on page 30 illustrates a simple loop connected to the host system through the SBus adapter. An SSA cable runs from SSA port 1 on the SBus adapter to a port on one end of the storage subsystem, and another SSA cable runs from a port on the other end of the subsystem to SSA port 2 on the SBus adapter, completing the loop.

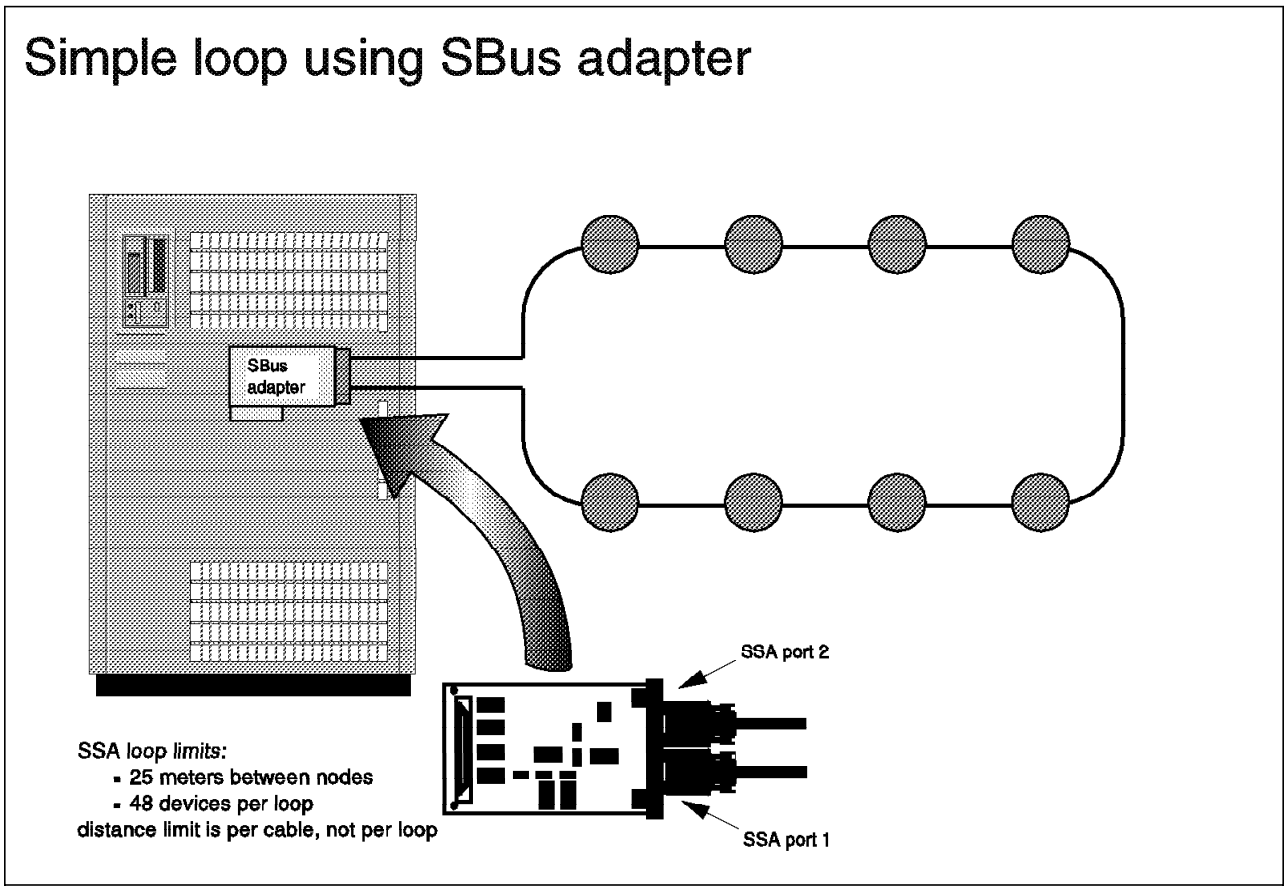

Figure 20. Simple Loop Using SBus Adapter.

Figure 21 on page 31 shows an equivalent configuration using the 7190. The subsystem is connected in the same fashion, with a cable from 7190 port A1 to the subsystem, and another cable from the subsystem to 7190 port A2. The only fundamental difference between the two configurations is the addition of a SCSI cable to connect the 7190 SCSI port to a fast wide differential SCSI adapter on the host system.

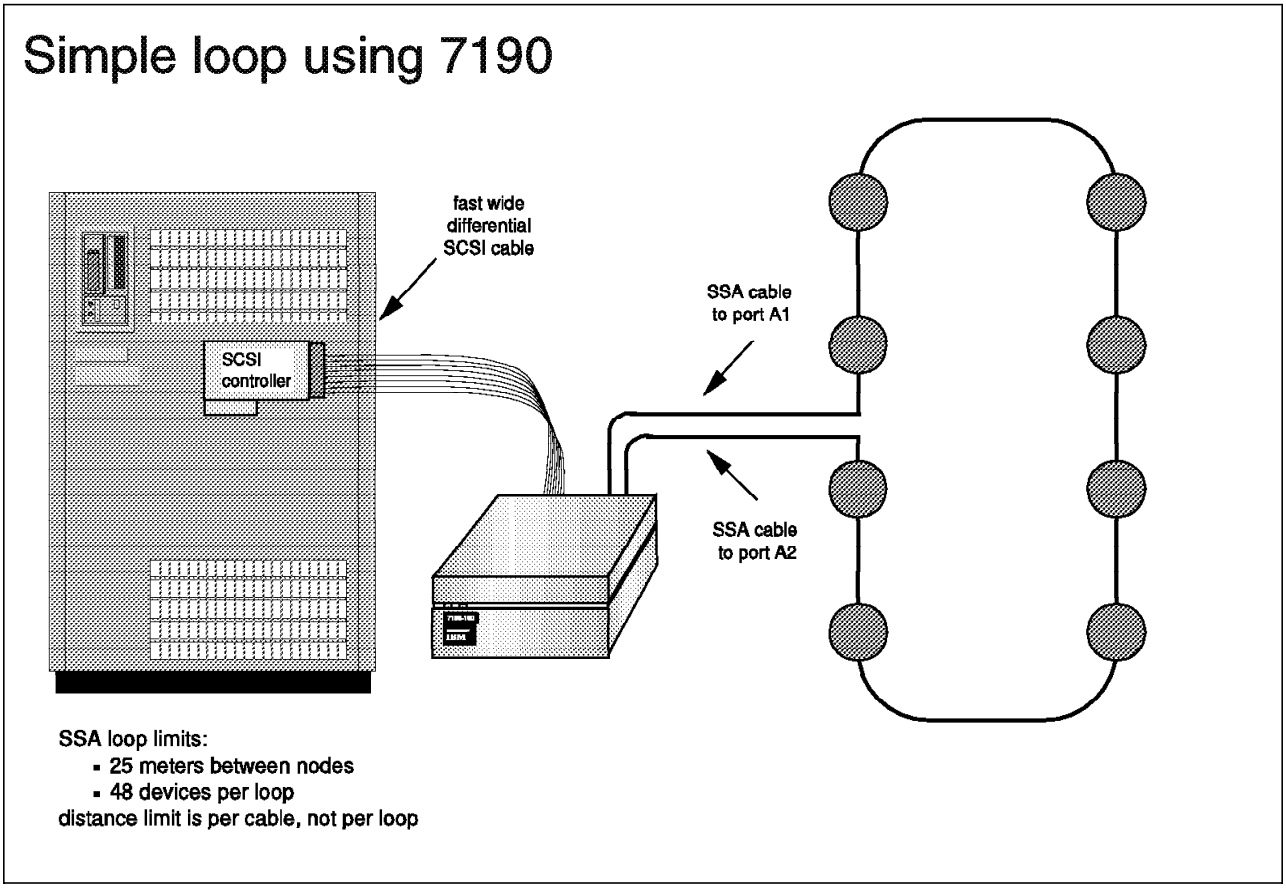

Figure 21. Simple Loop Using 7190.

3.3.2 One Enclosure or Several

The storage subsystems involved can be single enclosures with a few disk drives (see Figure 22) or several enclosures containing many disk drives and connected to each other with additional SSA cables (see Figure 23 on page 32). The cables used in both these configurations can be any of the available lengths of 1, 2.5, 5, 10, or 25 m each. An attribute of SSA that simplifies connectivity is the fact that the 25 m limit restricts only the distance between nodes, not the aggregate length of a string or loop.

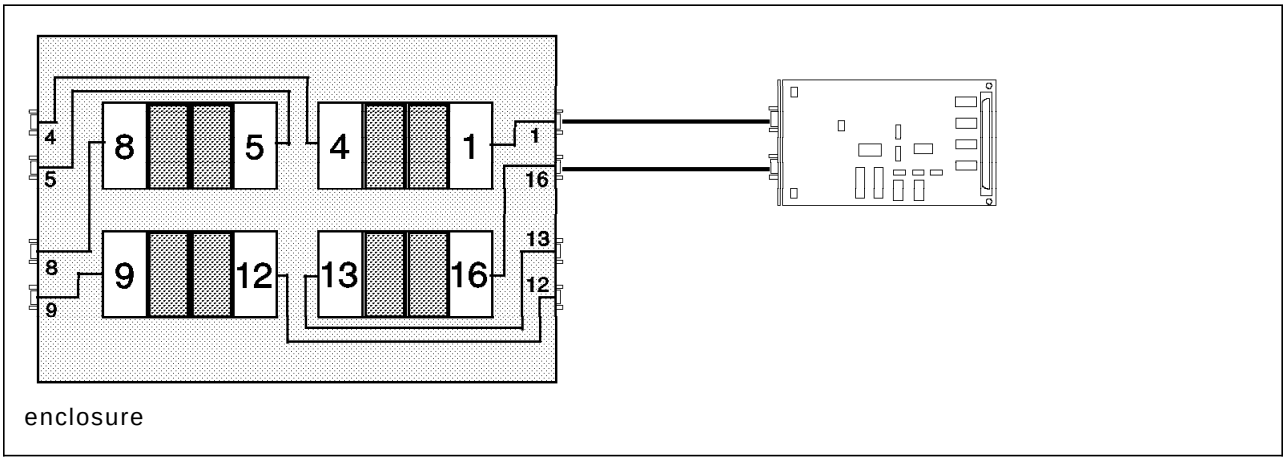

Figure 22. Single Enclosure in a Loop, Showing Eight Drives in One Enclosure

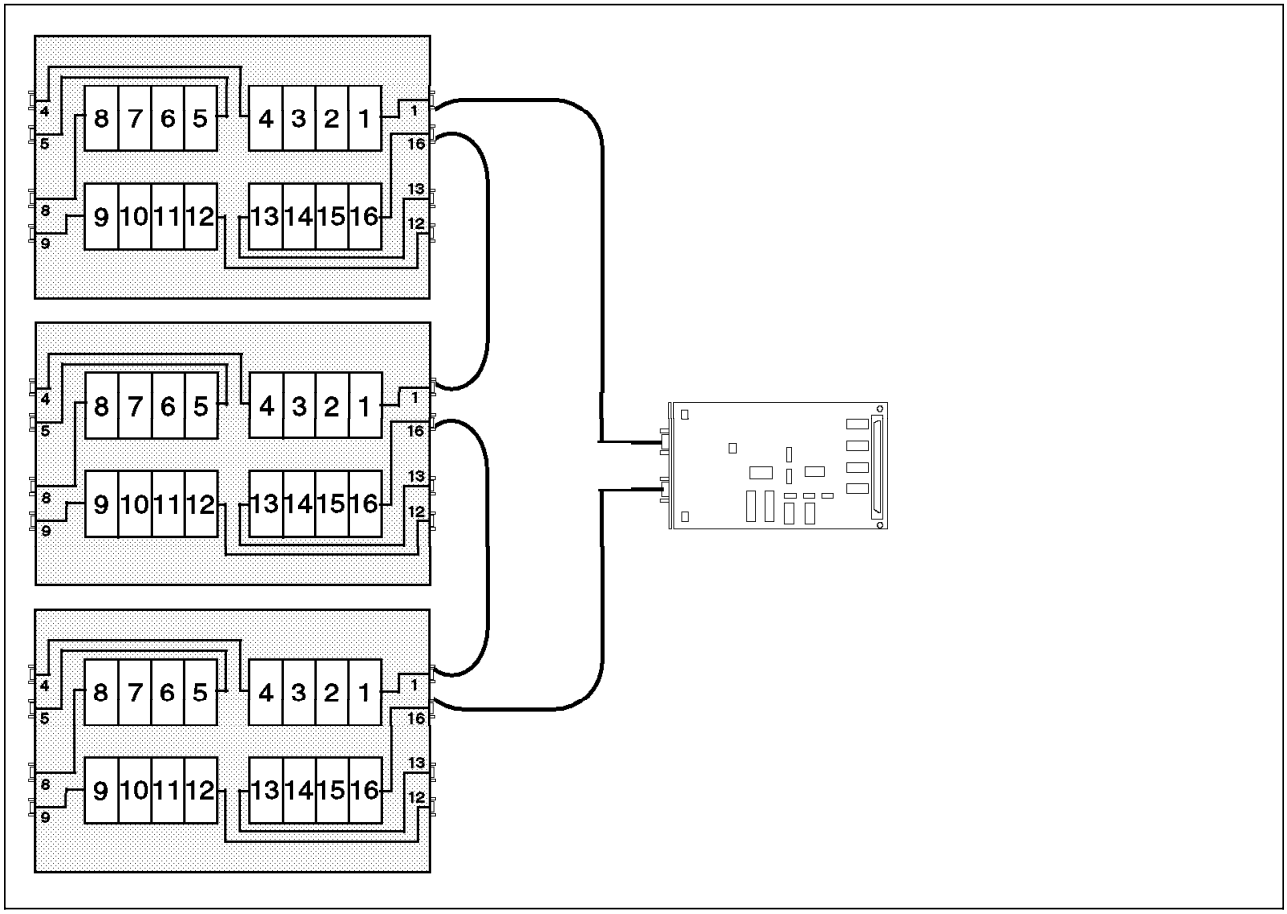

Figure 23. Multiple Enclosures in a Loop, Showing 16 Drives in Three Enclosures

3.3.3 Comparison with SCSI

For comparison, let's examine an equivalent SCSI configuration. An external SCSI cable is used to connect the host SCSI adapter to the subsystem. If the subsystem is made up of multiple enclosures, additional SCSI cables are used between them. Figure 24 on page 33 shows a sample of this configuration. Looking at distance restrictions, the best-case scenario in SCSI allows up to 25 m of aggregate bus length, using differential devices and cabling. Single-ended devices and cables reduce this to no better than 6 m. Compare this with the SSA distance of 25 m between nodes. For reference, Figure 25 on page 34 shows the length limits of various SCSI options.

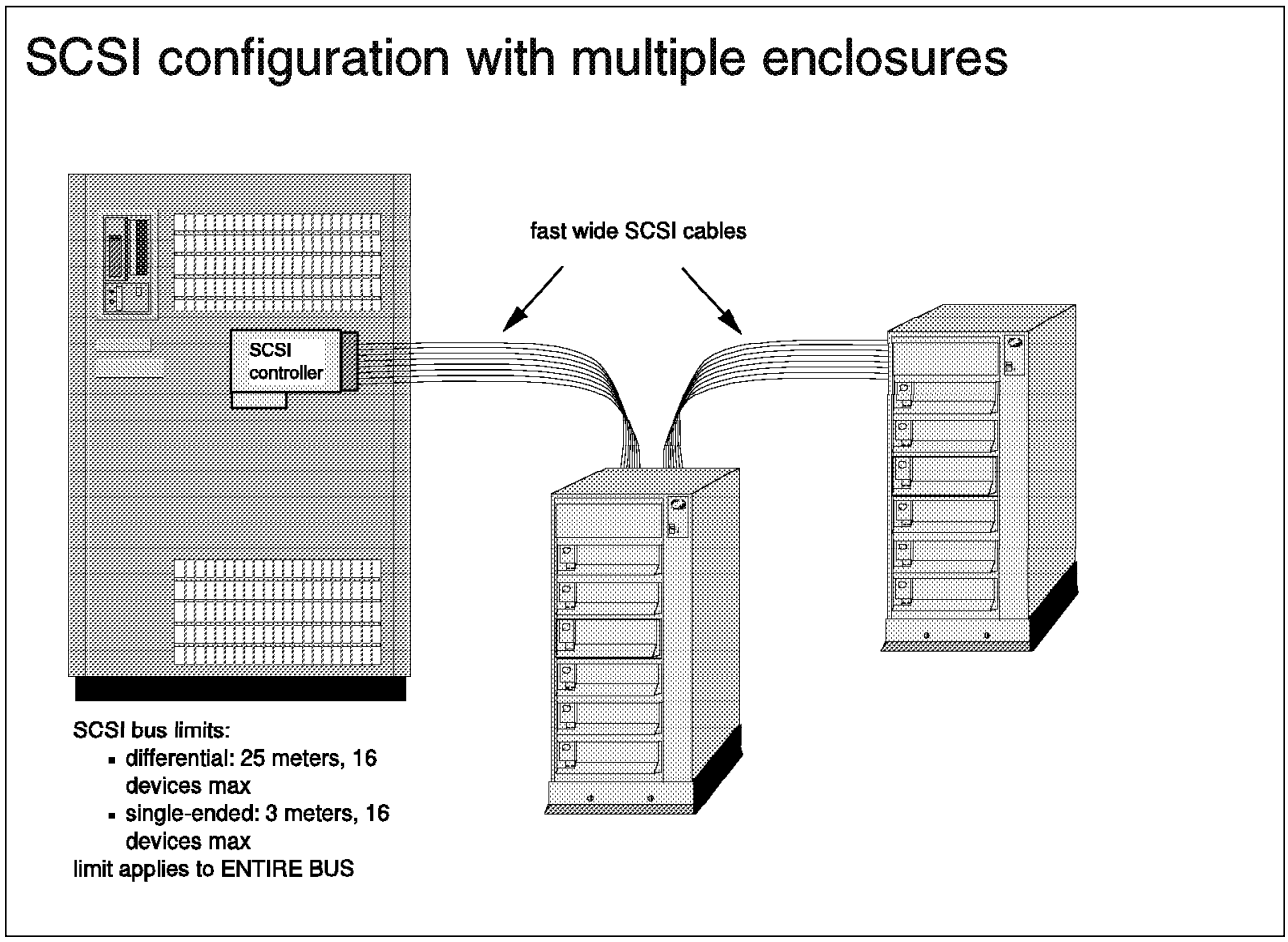

Figure 24. Sample SCSI Configuration.

Interface Type	Bus Speed, MB/s	Bus Width, Bits	Max Bus Length, Meters			Max Device Support
			Single-Ended	Differ-ential	LVD	
SCSI-1	5	8	6	25	12	8
Fast SCSI	10	8	3	25	12	8
Fast Wide SCSI	20	16	3	25	12	16
Ultra SCSI	20	8	1.5	25	12	8
Ultra SCSI	20	8	3	25	12	4
Wide Ultra SCSI	40	16	-	25	12	16
Wide Ultra SCSI	40	16	1.5	-	-	8
Wide Ultra SCSI	40	16	3	-	-	4
Ultra2 SCSI	40	8	*	*	12	8
Wide Ultra2 SCSI	80	16	*	*	12	16

*single-ended and differential not defined for Ultra2 SCSI

Figure 25. SCSI Line Lengths and Other Specifications for Various Implementations.

3.3.4 Fiber for Distance

Another connectivity option of SSA is the use of the fiber-optic extender (feature 5500 of the 7133). This device allows a greater distance limit of up to 2.4 km between SSA nodes. Figure 26 on page 35 shows an example of how this might be used (with a 7190, in this case). Port A1 of the 7190 is connected to the first enclosure in the loop using a standard SSA cable. This first enclosure is connected to a remote storage subsystem, located some distance away, using the fiber-optic extenders. The loop continues back to the first location, again using the fiber-optic extenders, and is completed with a standard SSA cable connecting the last enclosure to port A2 of the 7190. An SBus adapter could be used in exactly the same way. Figure 27 on page 36 shows a similar configuration with two hosts sharing storage from different locations within a campus environment.

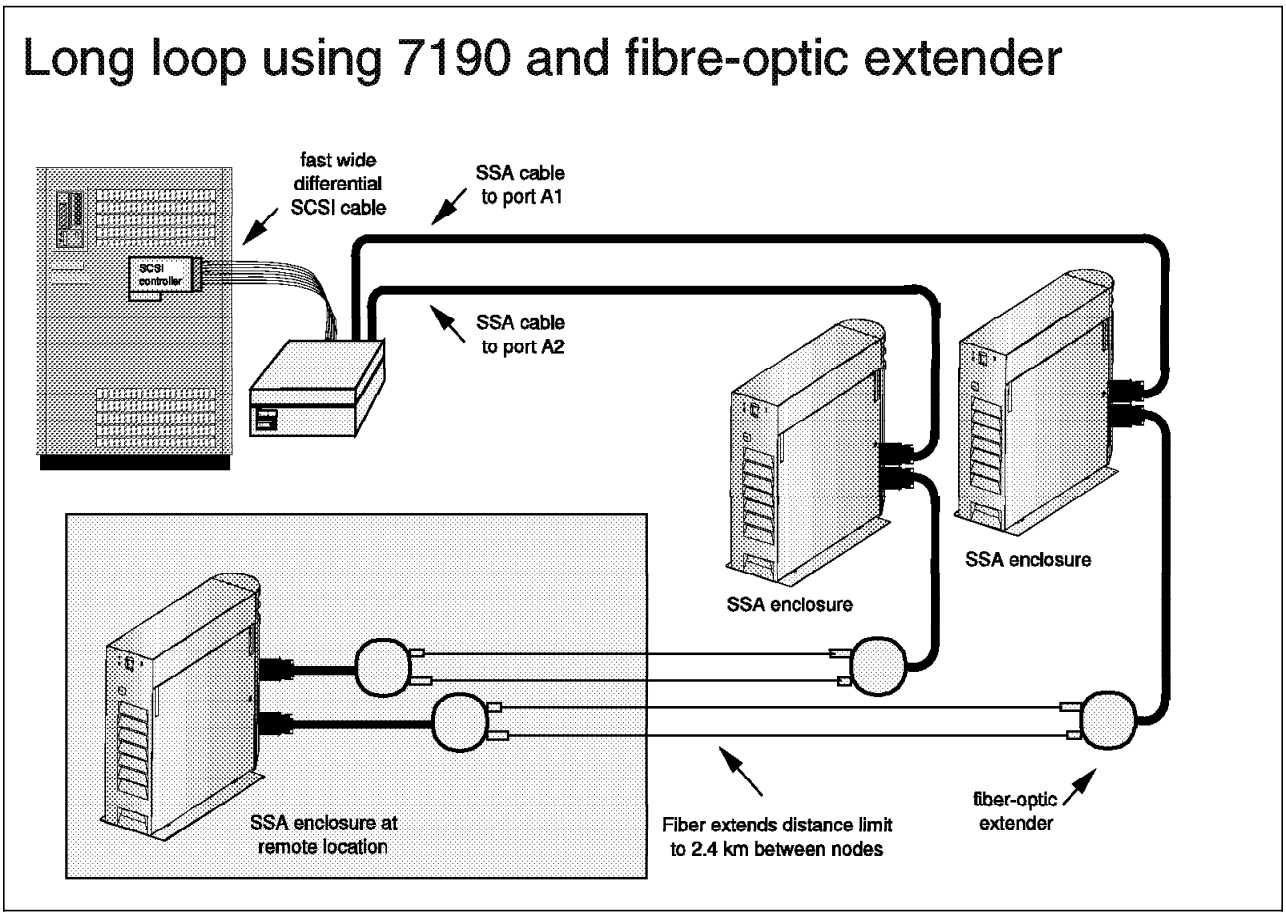

Figure 26. Use of Fiber-Optic Extender.

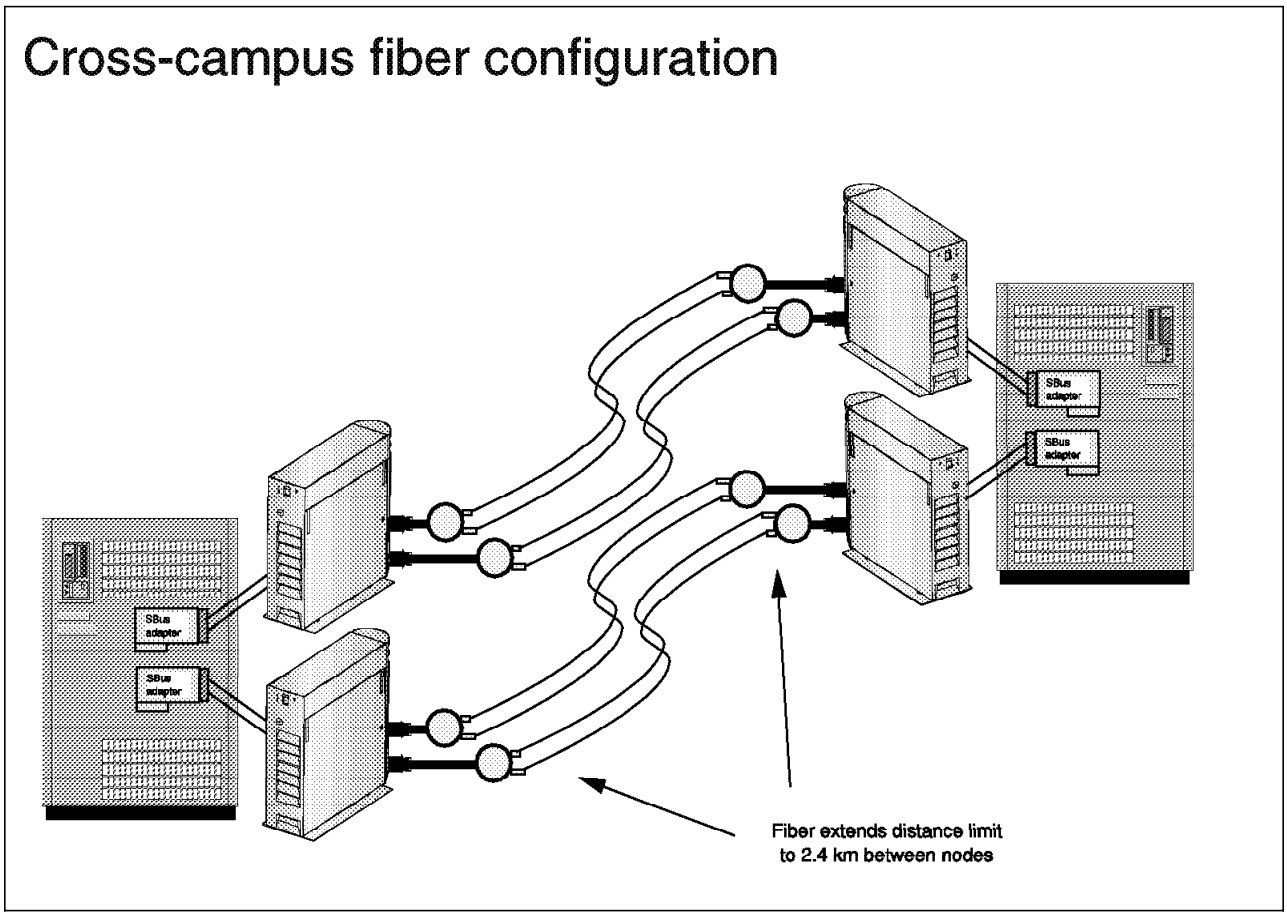

Figure 27. Multihost Use of Fiber-Optic Extender. This demonstrates use of the fiber-optic extender in a campus environment.

3.4 SSA Storage Subsystem Requirements

The 7190 and SBus adapter both allow connection to the entire family of IBM SSA storage subsystems. This family includes the IBM 7131 SSA multistorage Tower Model 405 (see Figure 28 on page 37), the IBM 7133 SSA subsystem Models 010 and 020 (see Figure 29 on page 38), and the IBM 7133 SSA subsystem Models 500 and 600 (see Figure 30 on page 39). This provides flexibility in selecting storage subsystems as it supports both new and mature SSA products.

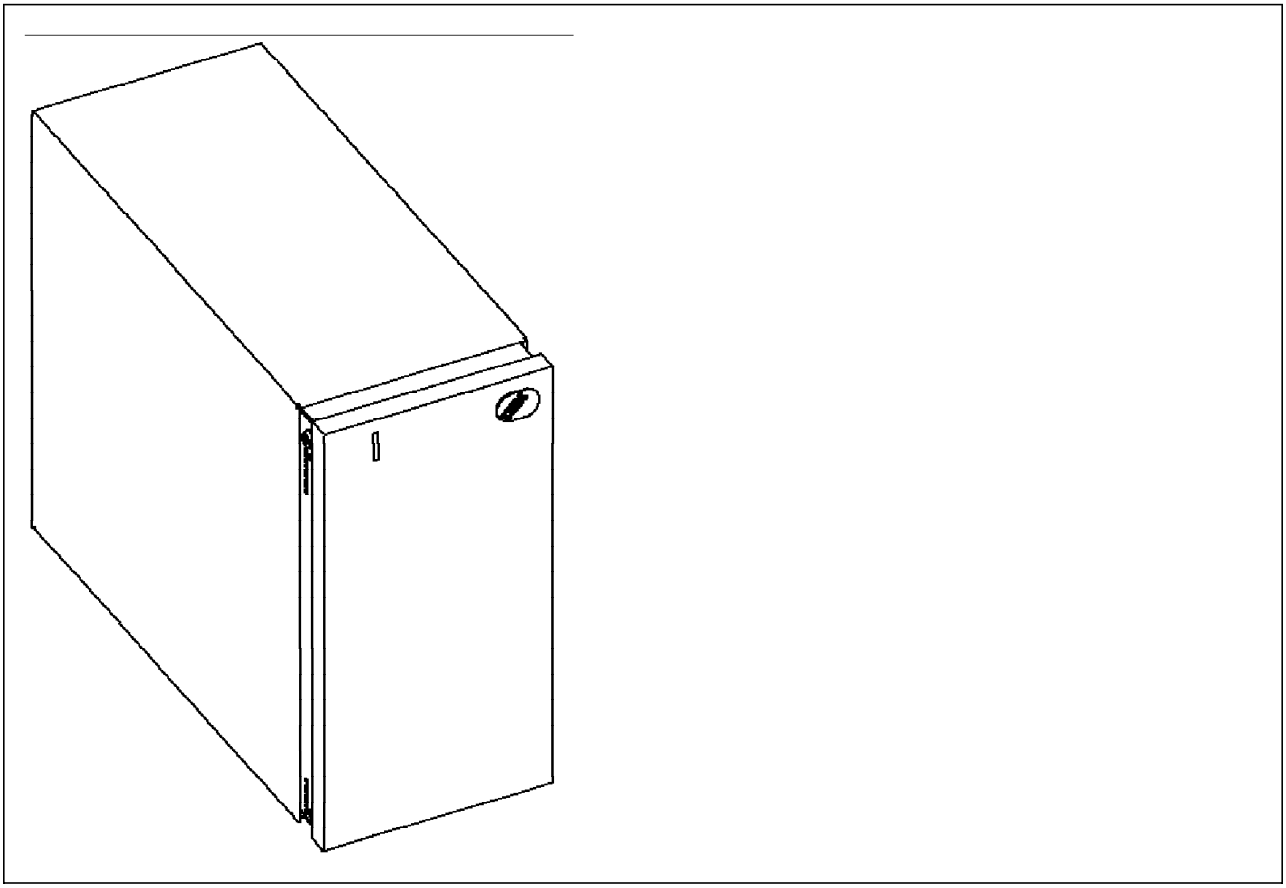

Figure 28. IBM 7131 SSA Multi storage Tower Model 405.

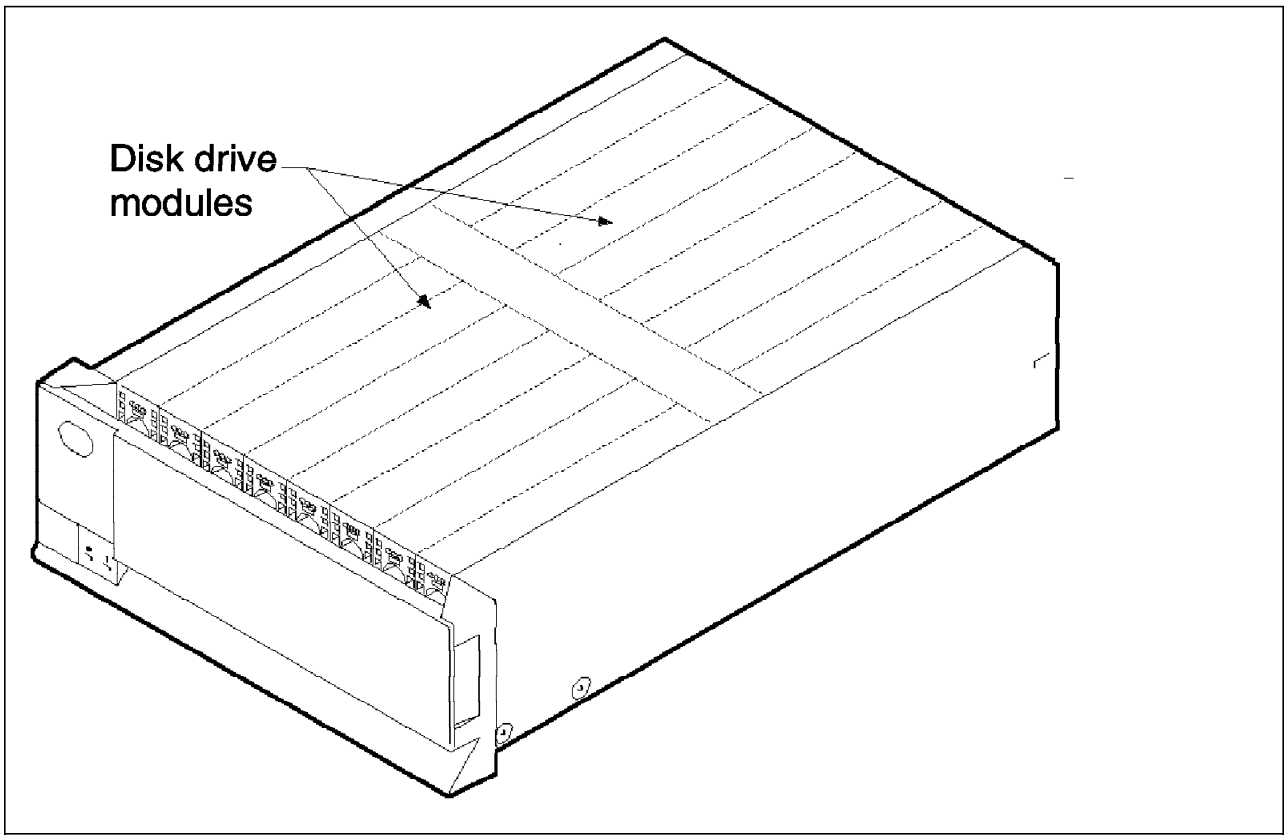

Figure 29. IBM 7133 SSA Subsystem Model 010/020.

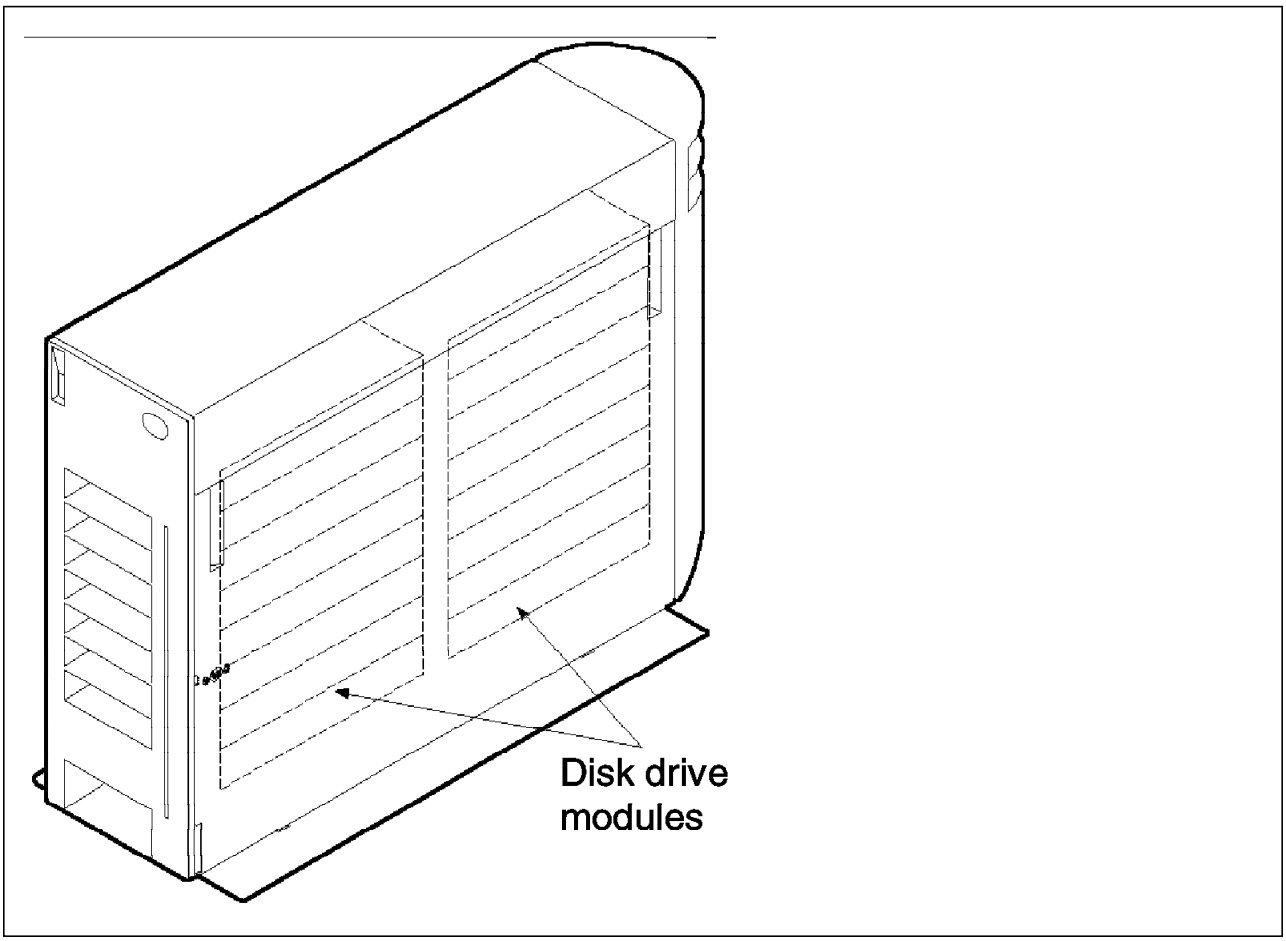

Figure 30. IBM 7133 SSA Subsystem Model 500/600.

When ordering the 7133 storage subsystem for attachment to the SBus adapter, specify 7133 Feature 4003. When ordering either the 7133 or 7131-405 storage subsystems for attachment to the 7190, specify 7133 or 7131-405 feature code 9901.

To attach an SSA storage subsystem to a Sun system through either the SBus adapter or the 7190, there are minimum firmware and microcode requirements that you must observe for both disk drives and the adapters. These firmware levels provide fixes for various issues encountered in product testing. Both products were rigorously tested using these levels, and product behavior using lower levels is unpredictable. Figure 31 lists the firmware requirements for the SBus adapter attachment. Figure 32 on page 40 lists the microcode requirements for 7190 attachment.

```
   SBus adapter firmware:                                1.50
   Ultrastar2 XP   Model DCHC 09 (9.1 GB) firmware:      dchc9292
   Ultrastar2 XP   Model DCHC 04 (4.5 GB) firmware:      dchc9290
   Ultrastar  XP   Model DFHC 04 (4.5 GB) firmware:      dfhc8877

Note: All of the above revisions are provided on the package 1.40 CD.
```

Figure 31. Minimum Disk and Adapter Firmware Requirements: SBus Adapter.

```
    7190 microcode:                                    6.8
    service function application:                      6.3
    Ultrastar2 XP   Model DCHC 09 (9.1 GB) microcode:  dchc9292
    Ultrastar2 XP   Model DCHC 04 (4.5 GB) microcode:  dchc9290
    Ultrastar  XP   Model DFHC 04 (4.5 GB) microcode:  dfhc8877

Note: The above microcode and service function revisions are provided on the
7190 Technical Support Home Page:
    http://www.storage.ibm.com/storage/techsup/7190/welcome.html
```

Figure 32. Minimum Disk and Adapter Microcode Requirements: 7190.

3.5 Design Considerations

When designing your storage subsystem, there are several topics that warrant careful consideration. In this section we address design issues specific to enhancing performance and availability. We discuss how to plan for and simplify future growth and point out some environment-specific considerations.

3.5.1 Designing for Performance

When designing your storage system to maximize performance, you can use several techniques, alone or in combination, depending on your specific application. In most cases, the techniques described here can be applied to either the 7190 or the SBus adapter.

3.5.1.1 Arrange Disks According to Utilization

Recall from 2.1.1.2, “Lower Layer Protocol” on page 10, that SSA frames are passed between nodes along a network. Each passage of a frame is called a *hop*. SSA packets are ultimately passed between a given disk drive and an adapter. The closer the drive is to the adapter, the fewer hops required for each frame to get from one to the other. It follows that minimizing the number of hops taken by the majority of SSA frames will improve performance because it reduces the amount of time spent by each node routing traffic for other nodes.

To use this characteristic to your advantage, arrange the disks along the loop according to expected utilization (or rearrange according to observed utilization). The disks with the highest utilization should be placed nearer the adapter, while those with lower expected utilization should be placed nearer the center of the loop (see Figure 33 on page 41). This arrangement will minimize the number of hops taken by the majority of the SSA frames, enhancing overall performance.

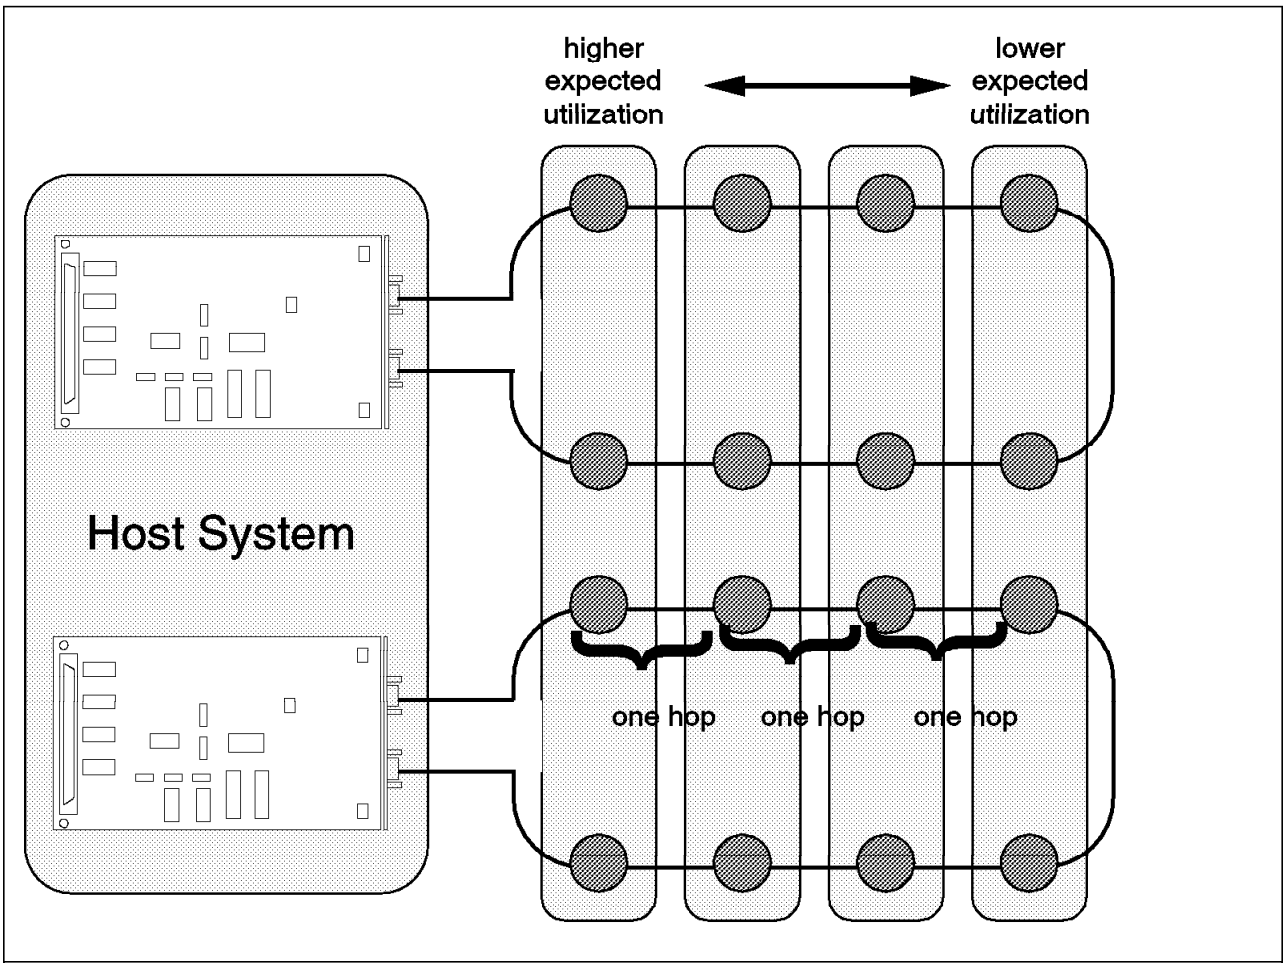

Figure 33. Arrange Disks According to Expected Utilization. Place the disks with highest expected utilization nearer the adapter.

3.5.1.2 Spread Workload across Adapters

If multiple adapters are to be used in the same host system, a performance enhancement can be realized by spreading the disks evenly between the adapters. One way of spreading the workload is to put the same quantity of disks on each adapter (see Figure 34 on page 42). This approach assumes that each of the disks has roughly the same I/O load. A slightly different approach can be applied if the I/O load varies significantly among disks. In this case, spread the disks between adapters in proportion to the amount of I/O activity expected on each (see Figure 35 on page 42). In either case, performance is maximized because the I/O activity is evenly shared between the two adapters, whether each disk has relatively equal I/O, or the adapters support equal I/O for an unequal number of disks. If this technique is used, it should certainly be combined with arranging disks along their respective loops according to utilization.

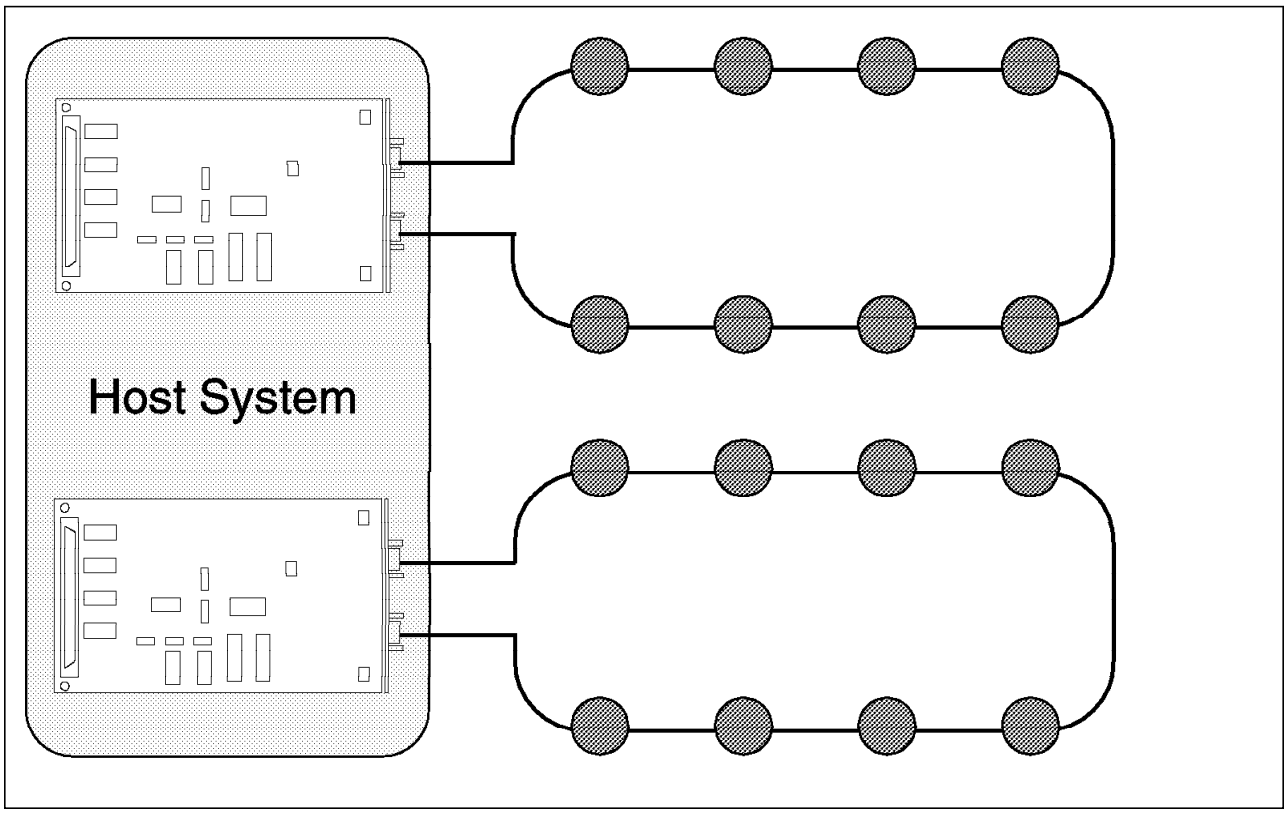

Figure 34. Number of Disks Spread Evenly among Adapters.

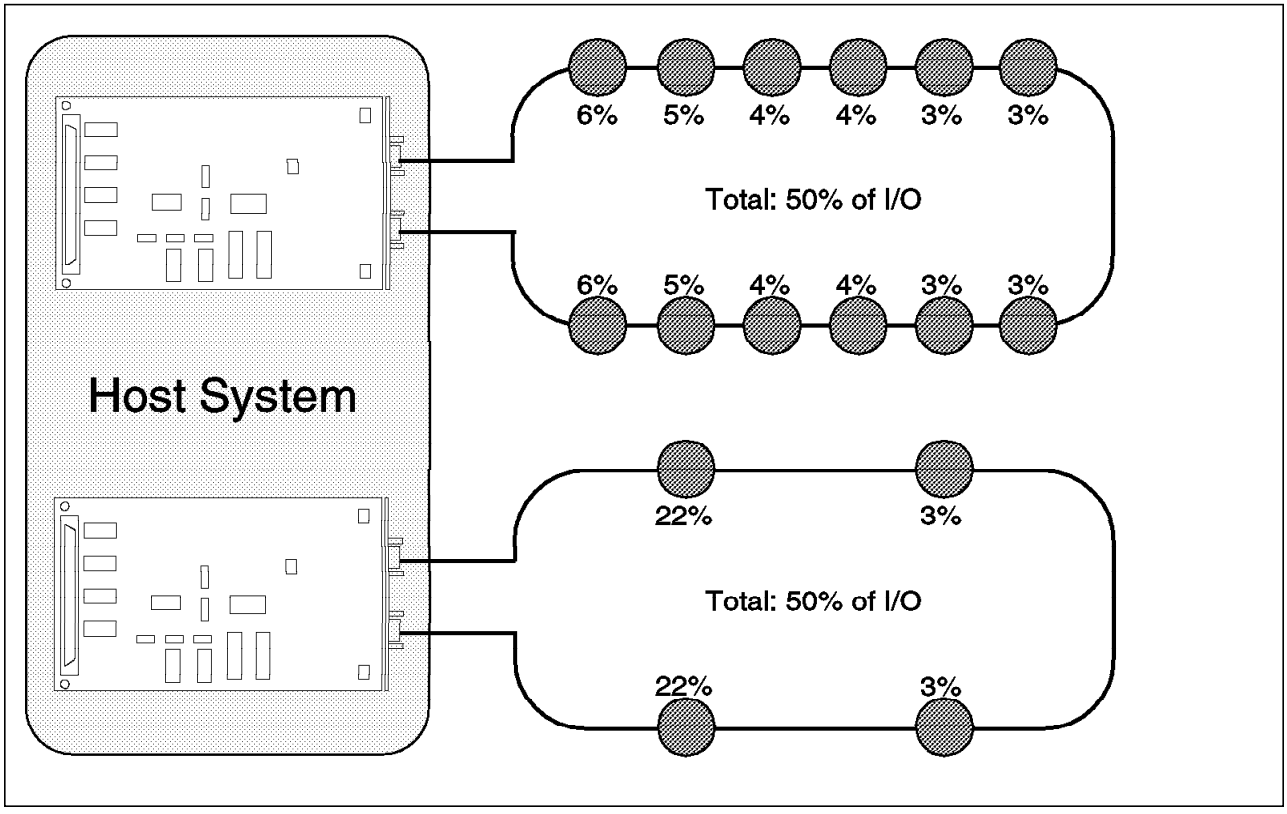

Figure 35. Disks Spread among Adapters in Proportion to Activity Level.

3.5.1.3 Place 7190 Alone on SCSI BUS

Bus arbitration is one of the characteristics of SCSI that interferes with performance. Only one device can transmit data or commands on a SCSI bus at a given time. Recall from Chapter 2, “SBus Adapter and 7190 Technology” on page 9, that SSA is not subject to this limitation. If a 7190 is located on a SCSI bus with other devices it will be subject to bus arbitration, impacting performance. To obtain maximum performance, we do not recommend installing the 7190 on a SCSI bus that includes other devices.

3.5.1.4 Implement RAID 0

The use of software striping, also known as RAID 0, can improve performance on the storage subsystem. In RAID 0, several disks are combined into a single volume, called a *stripeset*. Data within the volume is broken into small segments called *stripes*, which are spread evenly across the member disks (see Figure 36).

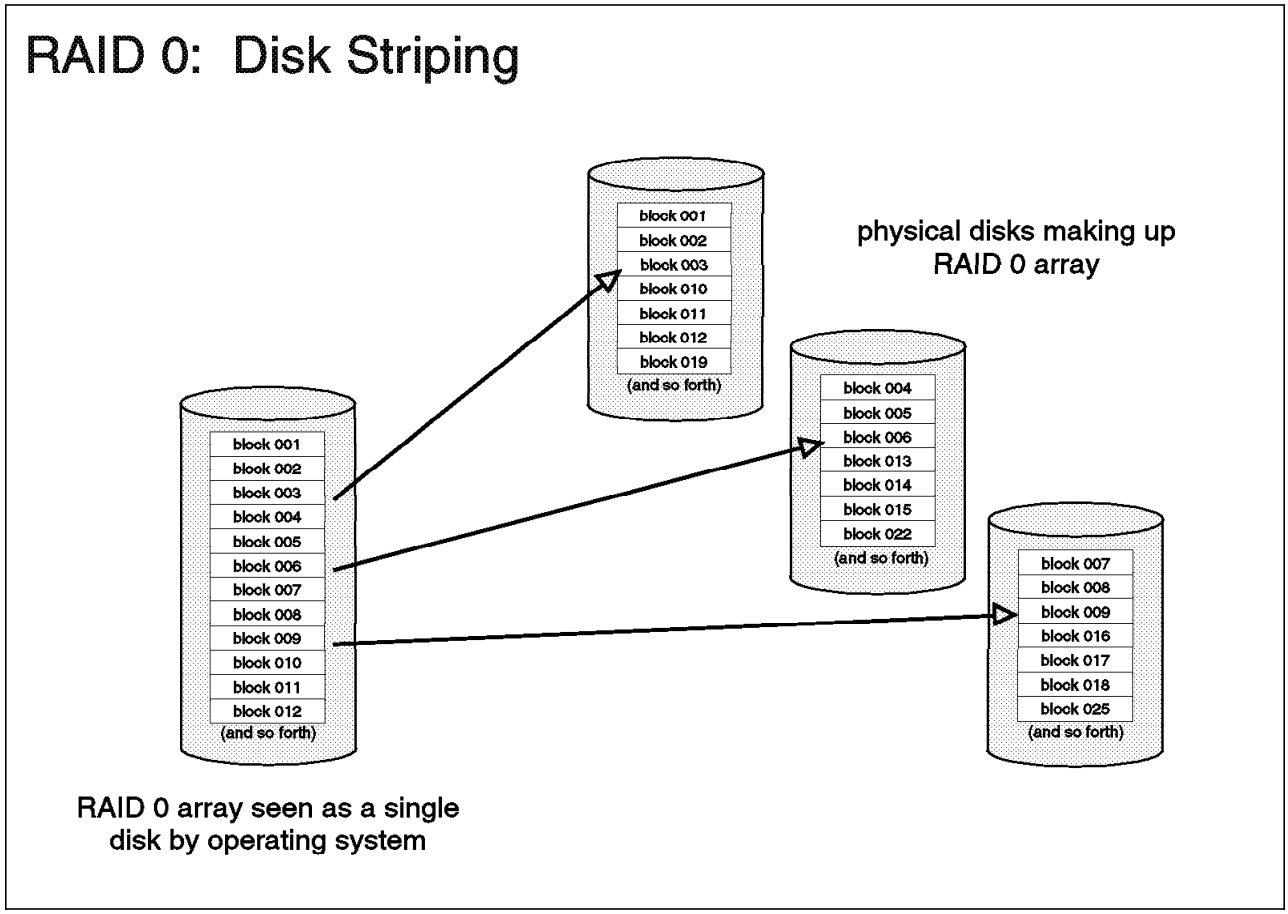

Figure 36. RAID 0: Striping.

RAID 0 provides performance gains over the same number of single disks by spreading the I/O across all members of the volume, rather than isolating the I/O to a single disk. This enhances performance in two ways: by balancing the request rate and by balancing the large-file transfer rate. For transaction requests where the request size is smaller than the stripe size of the array, a single disk retrieves or accepts the data for the request. In this situation, multiple requests can be satisfied concurrently by the array, improving the request rate over that of single disks (see Figure 37 on page 44). For large file requests where the request size is larger than the stripe size of the array, the

disks work in parallel to retrieve or accept the data for the request, improving transfer rate over that of single disks (see Figure 38 on page 45).

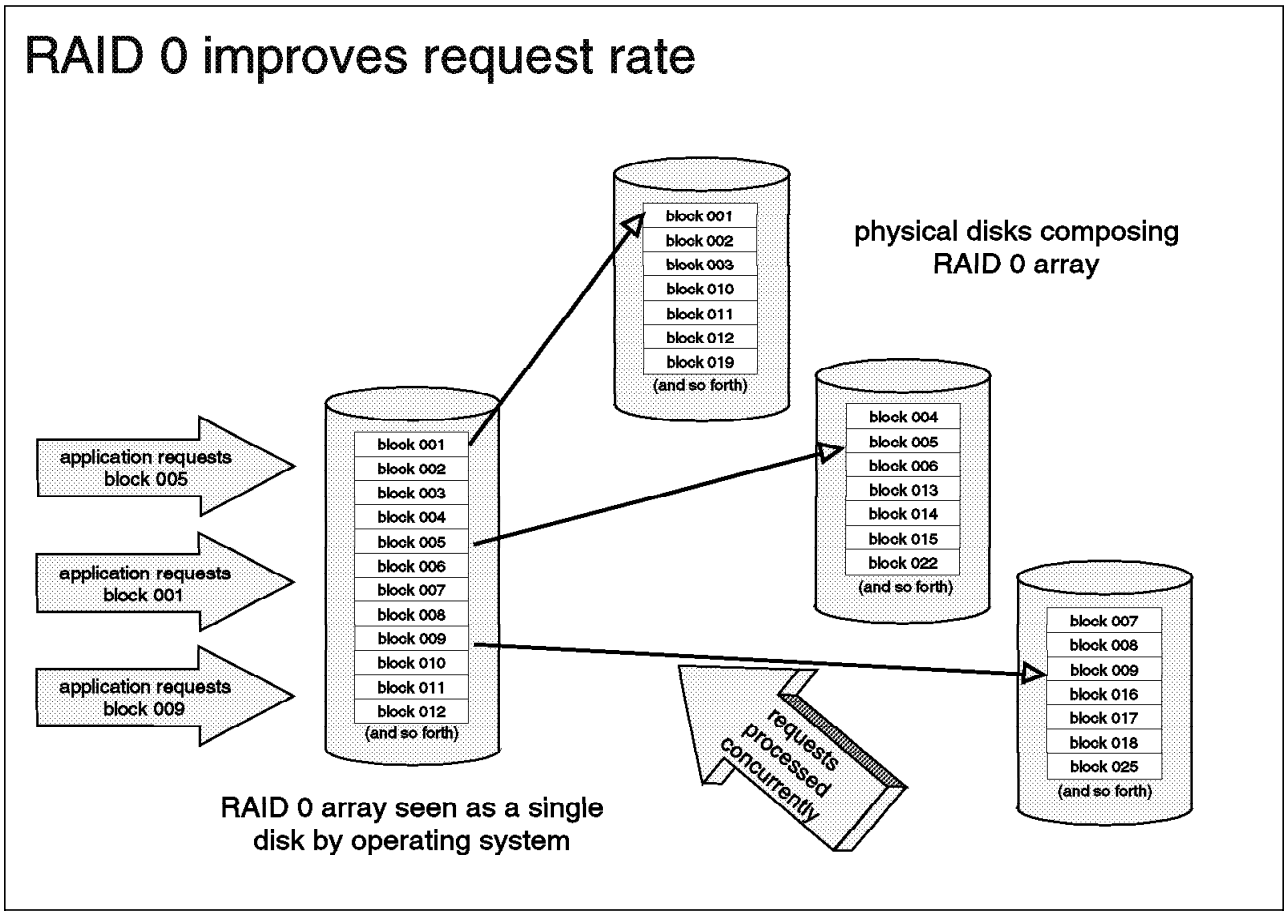

Figure 37. RAID 0 Improves Request Rate over That of Single Disks.

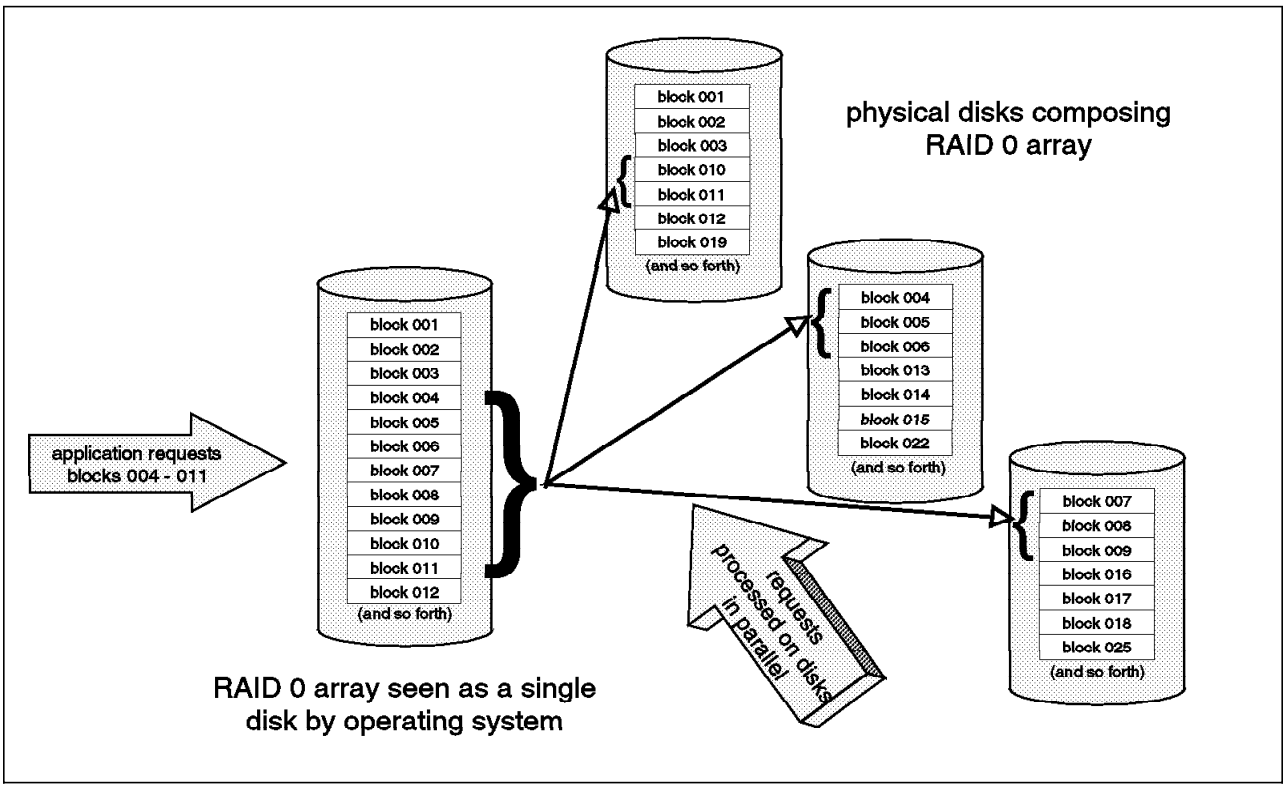

Figure 38. RAID 0 Improves Data Transfer Rate over That of Single Disks.

Several products support RAID 0 on the Solaris operating system, including Solstice DiskSuite and Veritas Volume Manager. When configuring an SSA subsystem using these products, the procedure you use is the same as for single SCSI disks. A single loop of SSA disks can be combined into one or several RAID 0 volumes. Several loops can also be used in a single volume.

RAID 0 performance can be further enhanced by spreading the members of the RAID set across multiple adapters. In this way, a given I/O workload is not only distributed among the stripeset members but also across the multiple adapters. In this case, both the disk seeks and the adapter I/O requests involved are balanced (see Figure 39 on page 46).

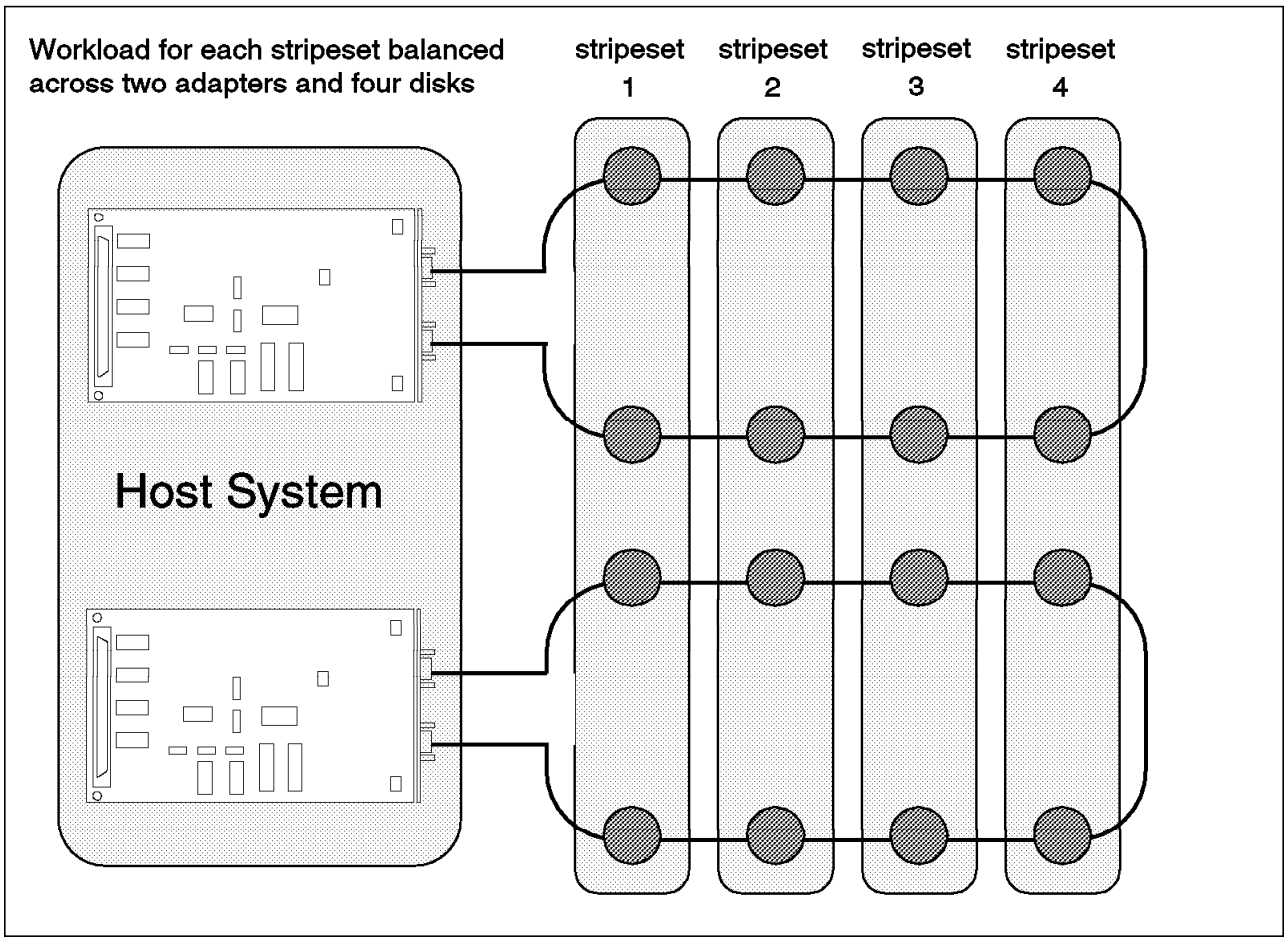

Figure 39. Balancing RAID 0 Seeks and Requests across Adapters.

Although RAID 0 provides a performance advantage, it is not without cost. For a given set of disks, RAID 0 increases the risk of data loss in terms of the quantity of data lost should a disk failure occur. The probability of disk failure does not change, but the quantity of data lost is greater if the failure does occur because the data on the entire RAID 0 stripeset is lost. You must evaluate this risk in light of the performance gains required on your particular subsystem.

3.5.2 Designing for Availability

When designing your storage system to maximize availability, your primary goal should be to identify and minimize the single points of failure within the planned system. How would the failure of a given component affect your overall availability of data? How would such a failure impact the ability of your applications to continue running? What impact would repair or routine maintenance of components have on the overall system? Finding the answers to these questions will help you design a system that experiences minimal or no impact from equipment failures and maintenance.

3.5.2.1 Single Host, Multiple Paths to Data

One availability problem with storage subsystems occurs with the failure of a storage adapter or the attached cabling. Although such failures are uncommon, careful planning of your storage arrangement can make your system tolerant of them. A technique that helps is to provide multiple paths to the same storage devices from a single host. In the SSA environment, you would place more than

one adapter attached to the same host in the same SSA loop (see Figure 40 on page 47). If one adapter fails, the remaining adapter still has access to the disks on the loop.

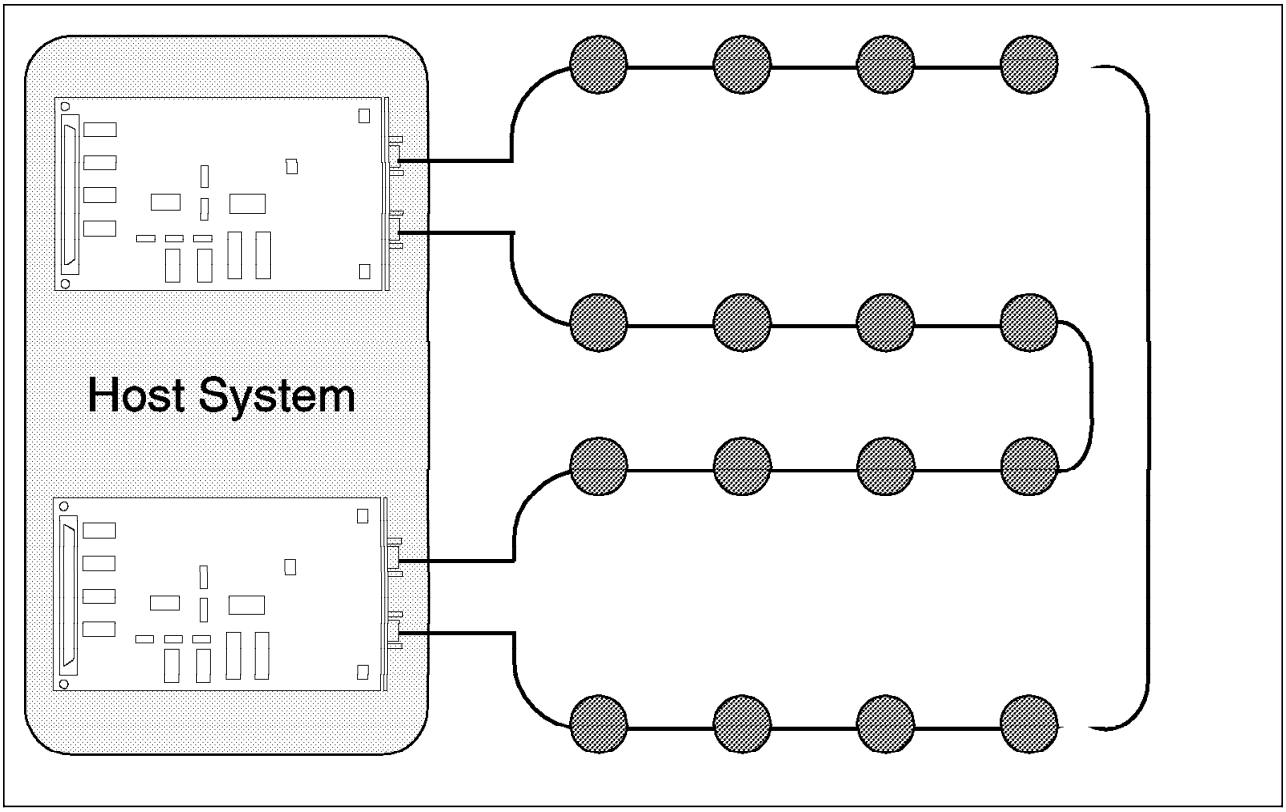

Figure 40. Single Host with Multiple Adapters. This configuration will maintain data access in the event of an adapter failure.

WARNING:

Using multiple paths to the same storage devices from the same host will present the operating system with two separate UNIX path names for the same physical disk, making it appear as two different disks (see Figure 41 on page 48). It is your responsibility as the user to know your configuration and which path names point to which devices. It is also your responsibility to configure your software environment so that both path names are not used at the same time. Failure to observe this warning can result in data corruption.

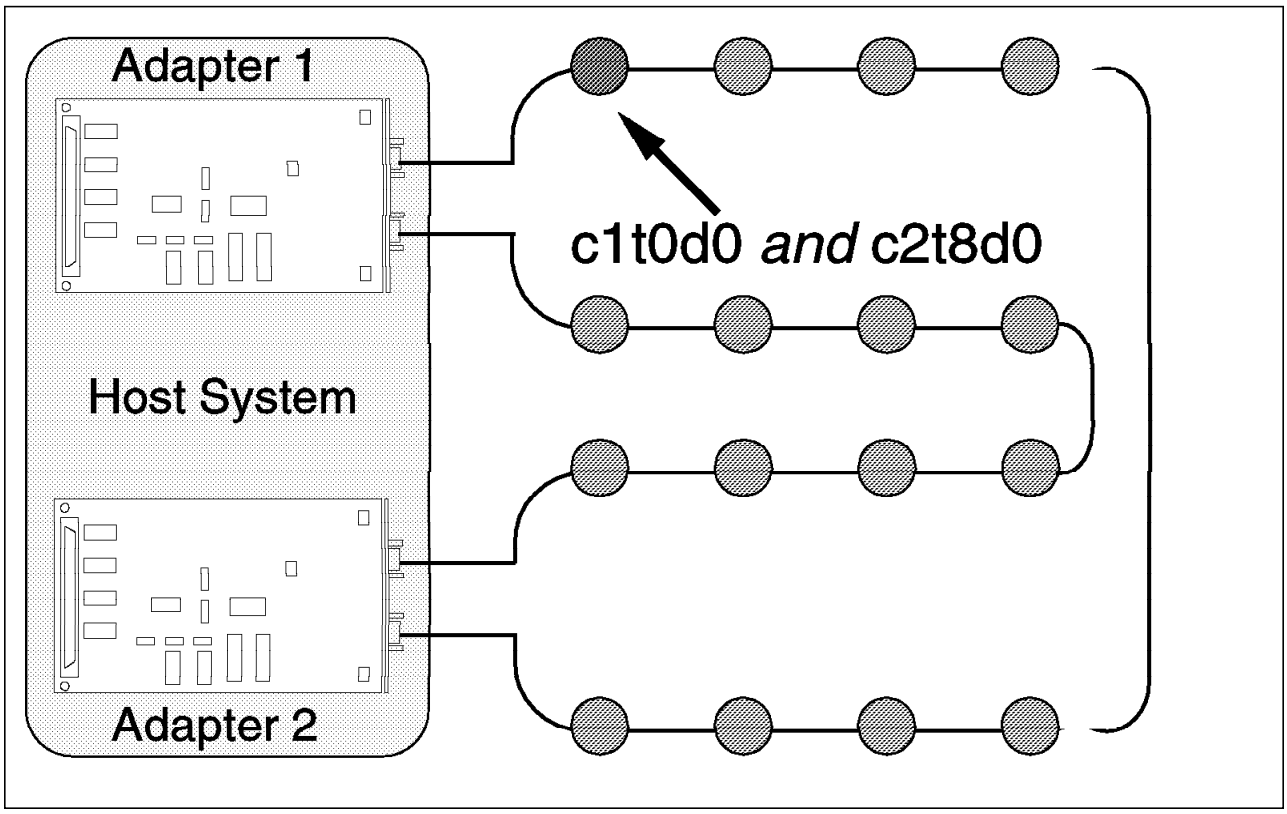

Figure 41. Two Path Names for the Same Disk. With two adapters in the same host connected to the same loop, the disk will appear to the host as two different disks: c1t0d0 and c2t8d0. Don't be fooled: both names actually apply to the same disk!

After giving due regard to the warning above, you can implement this configuration with either the 7190 or the SBus adapter, using two adapters on one host in the same loop.

3.5.2.2 Multiple Hosts, Single Path to Data

Another way to configure your storage system for availability is to use more than one host for data access. If one host loses access to the data, another can pick up this access. You configure multiple adapters in a loop, with each adapter in a different host (see Figure 42 on page 49). This configuration is supported for up to four adapters per loop by both the 7190 and the SBus adapter.

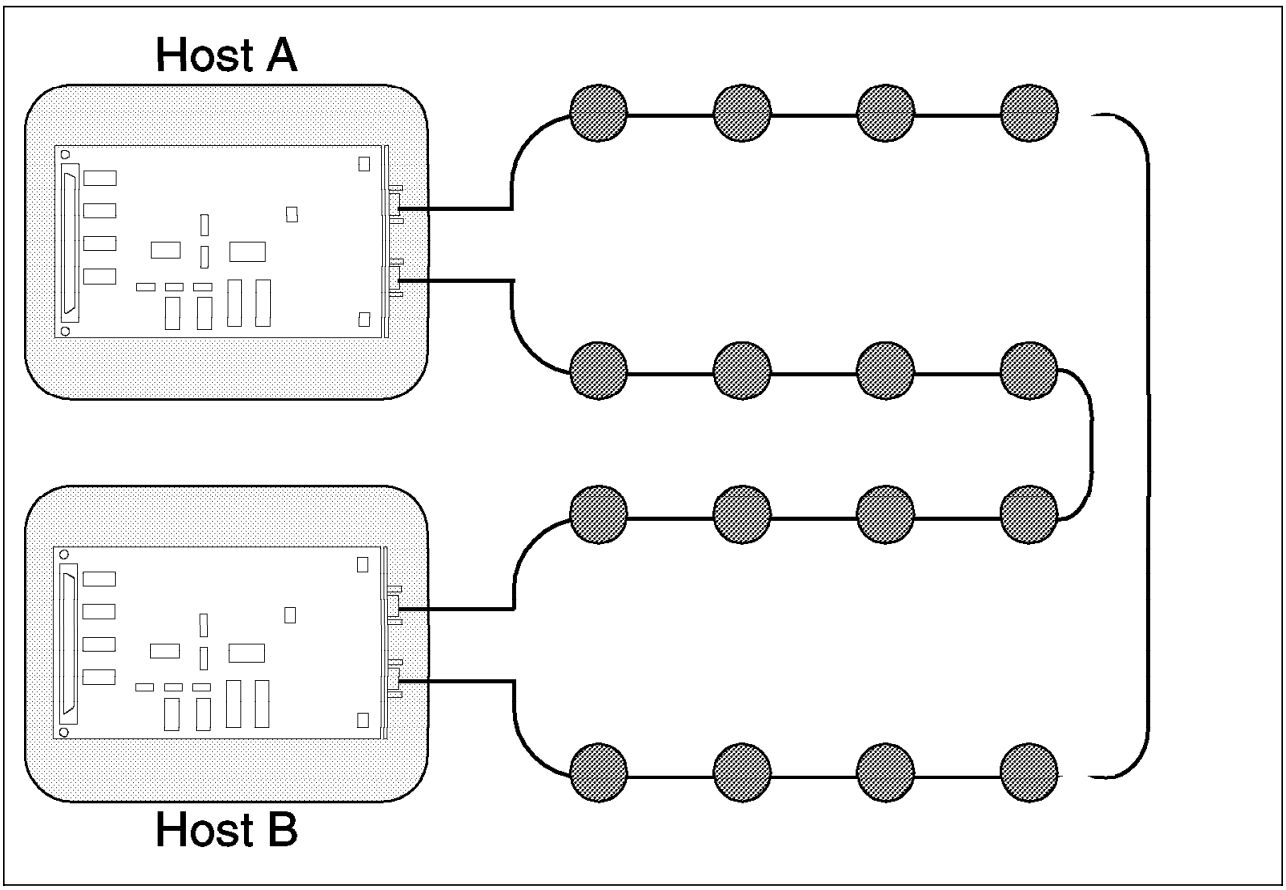

Figure 42. Multiple Hosts with Access to Same Loop. This configuration will maintain data access in the event of a host failure.

When configuring disks to be accessible by multiple hosts, access to the disks must be synchronized to ensure data integrity. If two hosts were to access the same disk at the same time, data corruption would probably occur. In the Solaris environment, applications such as Solstice HA, Oracle Parallel Server, and Veritas FirstWatch synchronize data access by ensuring that only one host has access to a given device at any one time.

3.5.2.3 Multiple Hosts, Multiple Paths to Data

The previous two high-availability configurations can be combined by setting up multiple hosts with multiple adapters on the same loop (see Figure 43 on page 50). This configuration, when configured with appropriate high availability software, provides tolerance of both a data path failure and a host system failure.

WARNING

Using multiple paths to the same storage devices from the same host will present the operating system with two separate path names for the same physical disk, making it appear as two different disks (see Figure 41 on page 48). It is your responsibility as the user to know your configuration and which path names point to which devices. It is also your responsibility to configure your software environment so that both path names are not used at the same time. Failure to observe this warning can result in data corruption.

After giving due regard to the warning above, you can implement this configuration with the either the 7190 or the SBus adapter, using up to two adapters per host per loop, with a maximum total of up to four adapters per loop.

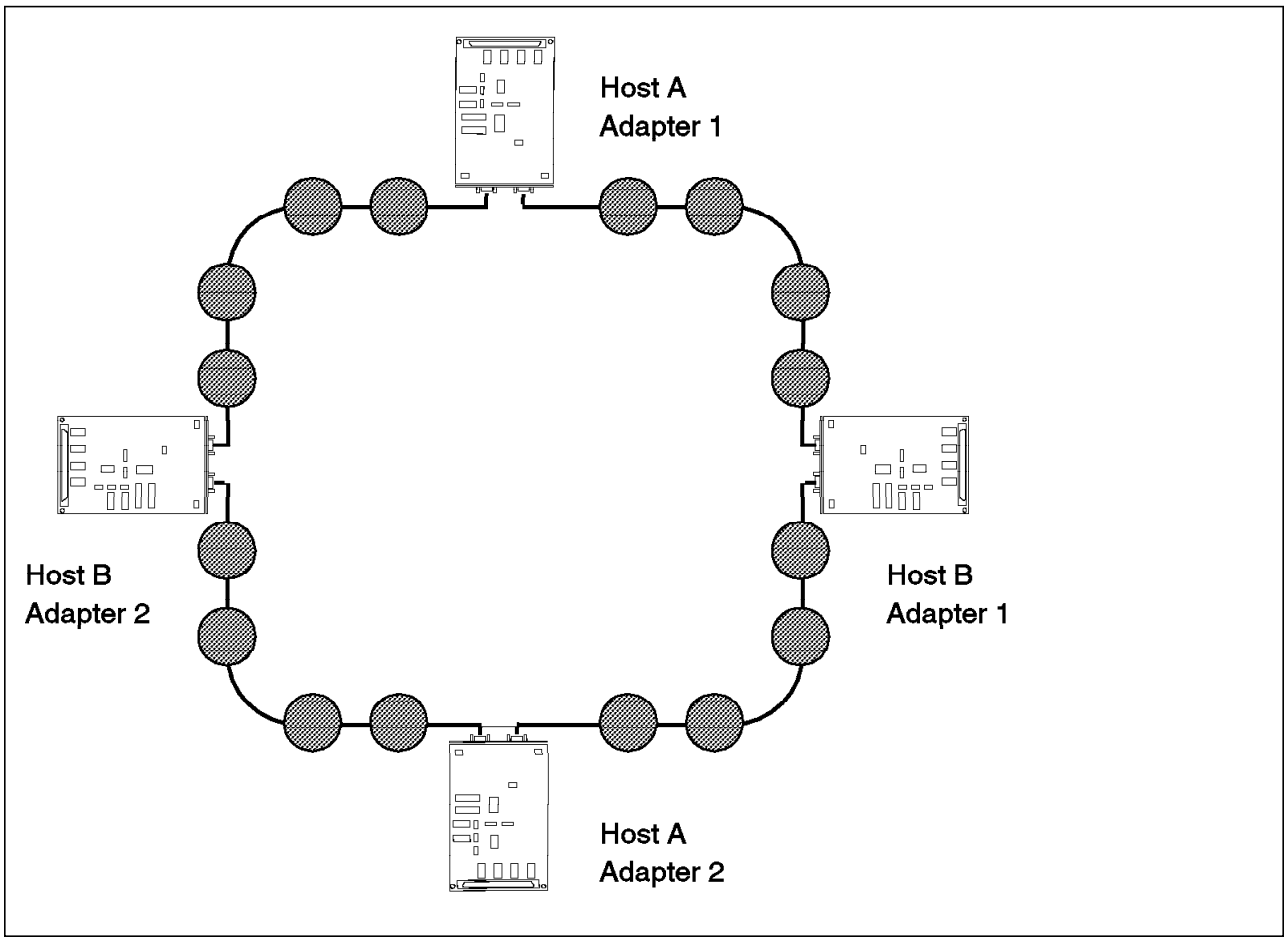

Figure 43. Multiple Hosts with Multiple Adapters on Same Loop. This configuration will maintain data access in the event of either a host failure or an adapter failure.

3.5.2.4 RAID for Availability

Another factor to consider when configuring a storage subsystem for availability is the use of software RAID, provided by products such as Solstice DiskSuite or Veritas Volume Manager. Three different RAID levels are commonly used for data availability purposes. Each allows a disk array to tolerate at least one disk failure and still continue operation and access to data. We cover these RAID levels briefly here.

RAID 5: Striping with Parity: RAID 5 makes a single volume, sometimes called a *parity stripeset*, from an array of three or more equally sized disks. The total capacity of the volume is the sum of the capacities of all but one of the disks (size * (n-1)). The disks are striped such that both data and parity are spread evenly across the disks.

RAID 1: Mirroring: RAID 1 presents the operating system with a single volume, called a *plex* or *mirror set*, that consists of two (or more) equally sized disks. For simplicity of discussion, we assume that all mirror sets consist of two members. The volume is seen as a single disk with a total capacity equal to that

of a single member. RAID 1 offers full redundancy of data because each mirror set member is an identical copy of the other. When all members are functioning, writes take place to all members. Depending on the implementation, reads can be spread among members, which improves performance over that of a single disk. If a member disk fails, I/O continues on the other member or members. The failure is transparent to the application.

RAID 0+1: Striped Mirroring: RAID 0+1 combines the data redundancy of mirroring with the performance of striping. RAID 1 mirror sets are created from member disks. Each mirror set is thus seen as if it were a single disk. These mirror sets are then combined into a RAID 0 stripeset. Performance is as good or better than that of RAID 0 alone, with the added protection that each member "disk" of the RAID 0 array is actually a RAID 1 mirror set. If a disk fails, I/O continues on the remaining members. In fact, up to half the disks can fail and the volume is still intact, provided no two members of the same RAID 1 mirror set have failed (assuming two-member RAID 1 sets).

Some implementations reverse the relationship of RAID 0+1, creating a RAID 1 mirror set out of two or more RAID 0 stripesets (mirrored stripesets), rather than creating a RAID 0 stripeset out of several RAID 1 mirror sets (striped mirror sets). See Figure 44 for an illustration of this.

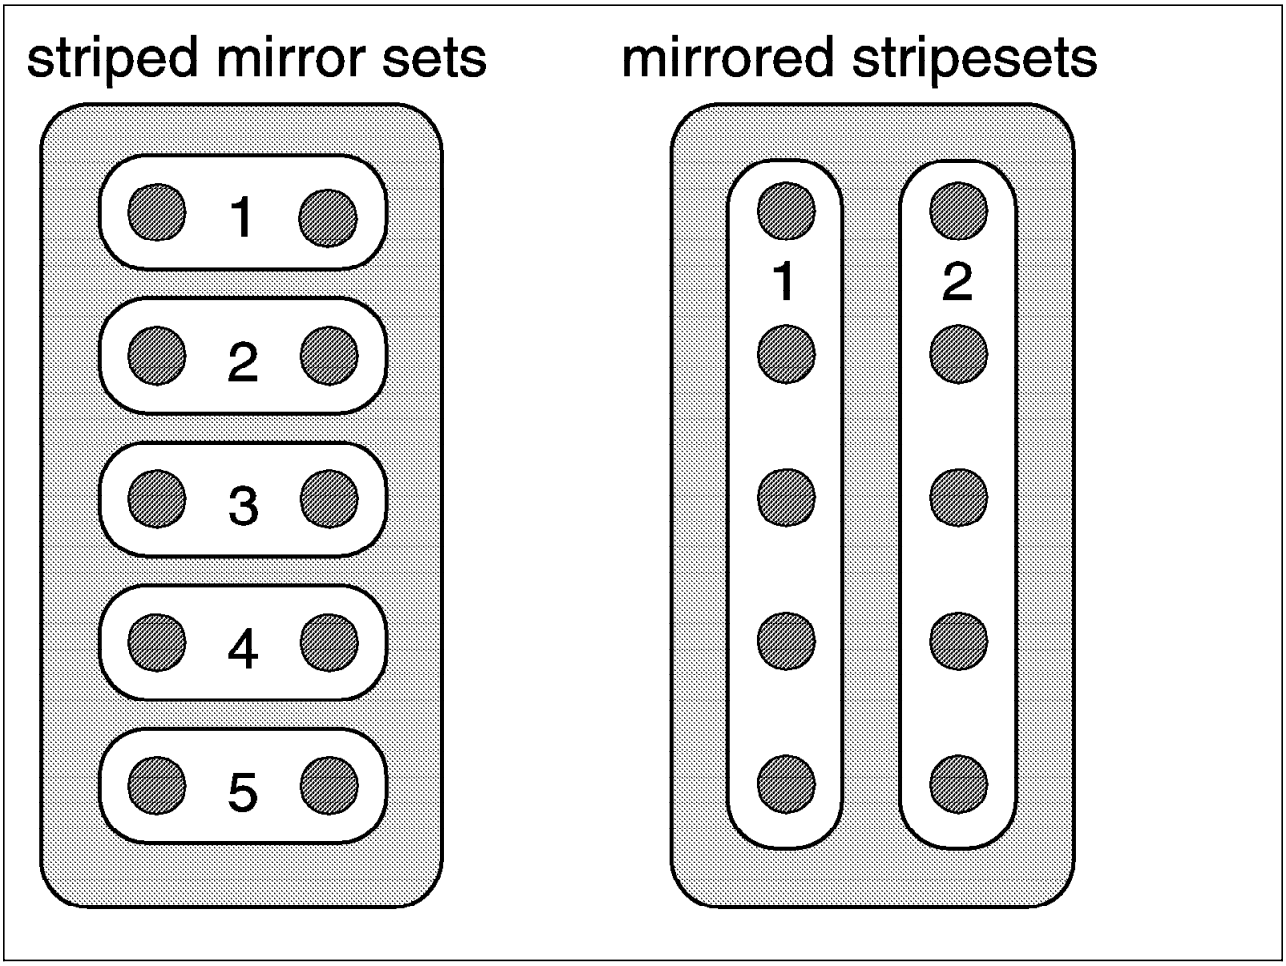

Figure 44. Two Different RAID 0+1 Implementations. The volume on the left consists of striped mirror sets, while the volume on the right consists of mirrored stripesets.

The subtle difference between this scenario and the previous scenario is one of availability. Both implementations provide full protection against a single disk failure. Of greater interest is the impact of a second disk failure, however unlikely that may be. In mirrored stripesets, losing a disk renders the entire stripeset invalid. At this point, no redundancy is provided by that mirror set member. Failure of any disk on the remaining stripeset renders that stripeset invalid, and you have lost the entire volume. With striped mirror sets, operations can continue with a loss of up to half the member drives, as long as no two members of the same mirror set have failed. This is illustrated in Figure 45.

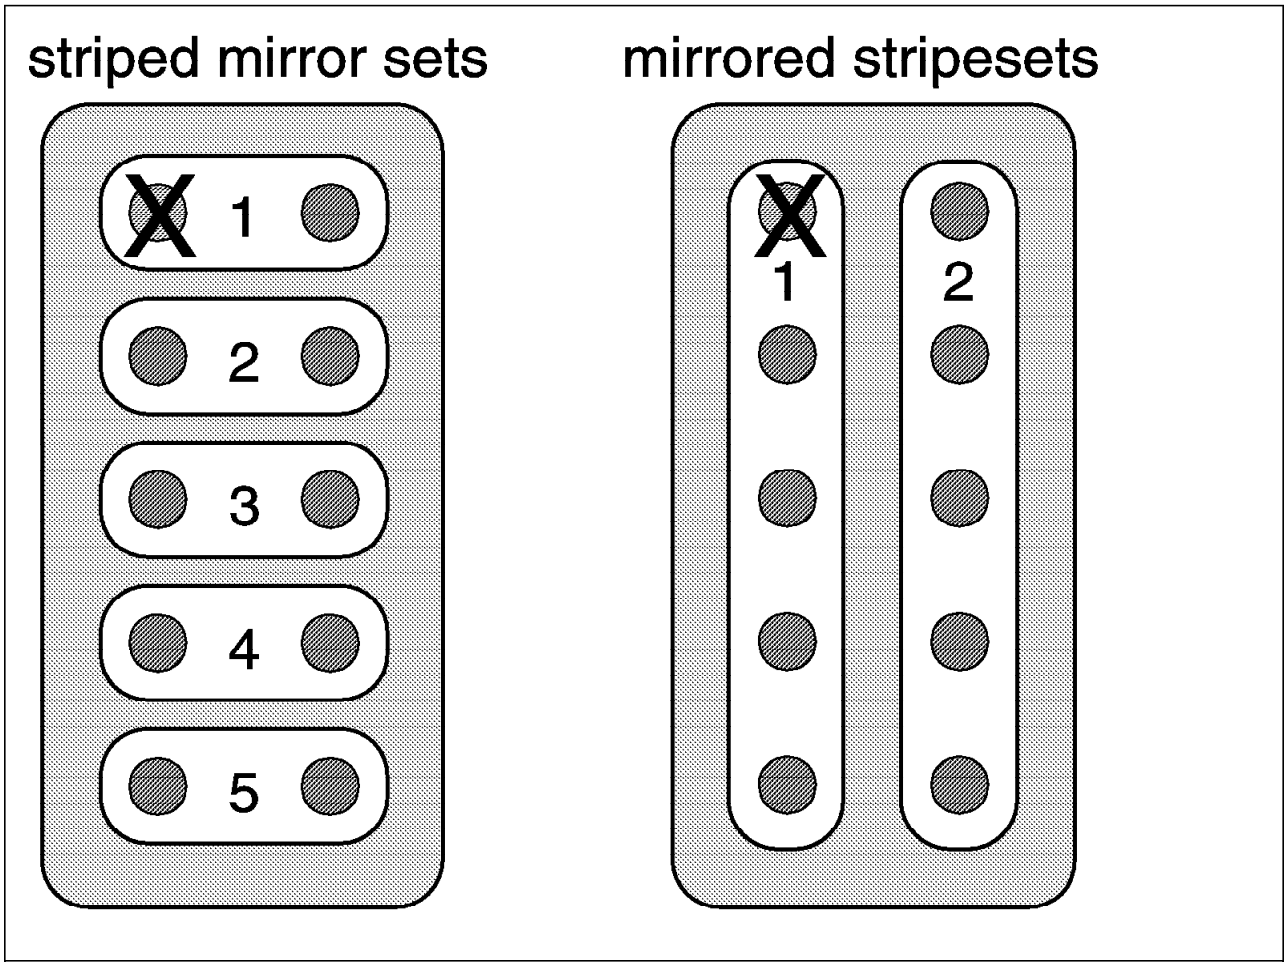

Figure 45. Impact of Disk Failure on RAID 0+1 Implementations. Both configurations have a failed disk. The configuration on the left has lost one member of mirror set 1, and can continue with the other member. This configuration will tolerate the failure of any disk except the second member of mirror set 1. The configuration on the right has lost all of stripeset 1, due to the failure of one member. If any member if stripeset 2 fails, the entire volume is lost.

3.5.2.5 Higher Availability: Mirroring across Adapters

Even though using RAID 1 or RAID 0+1 is a good way of preserving data availability, you can use other techniques to further enhance availability and realize a performance gain as well. You can arrange mirror sets so that members of each set are on different adapters, as shown in Figure 46 on page 53. Because I/O to the array is now spread across these adapters, performance is better than that of a single adapter. Availability is improved because the mirror set can survive the loss of an entire adapter. Arrange the

arrays so that members of each mirror set are in separate enclosures with independent power, and you have built-in tolerance of a power supply failure.

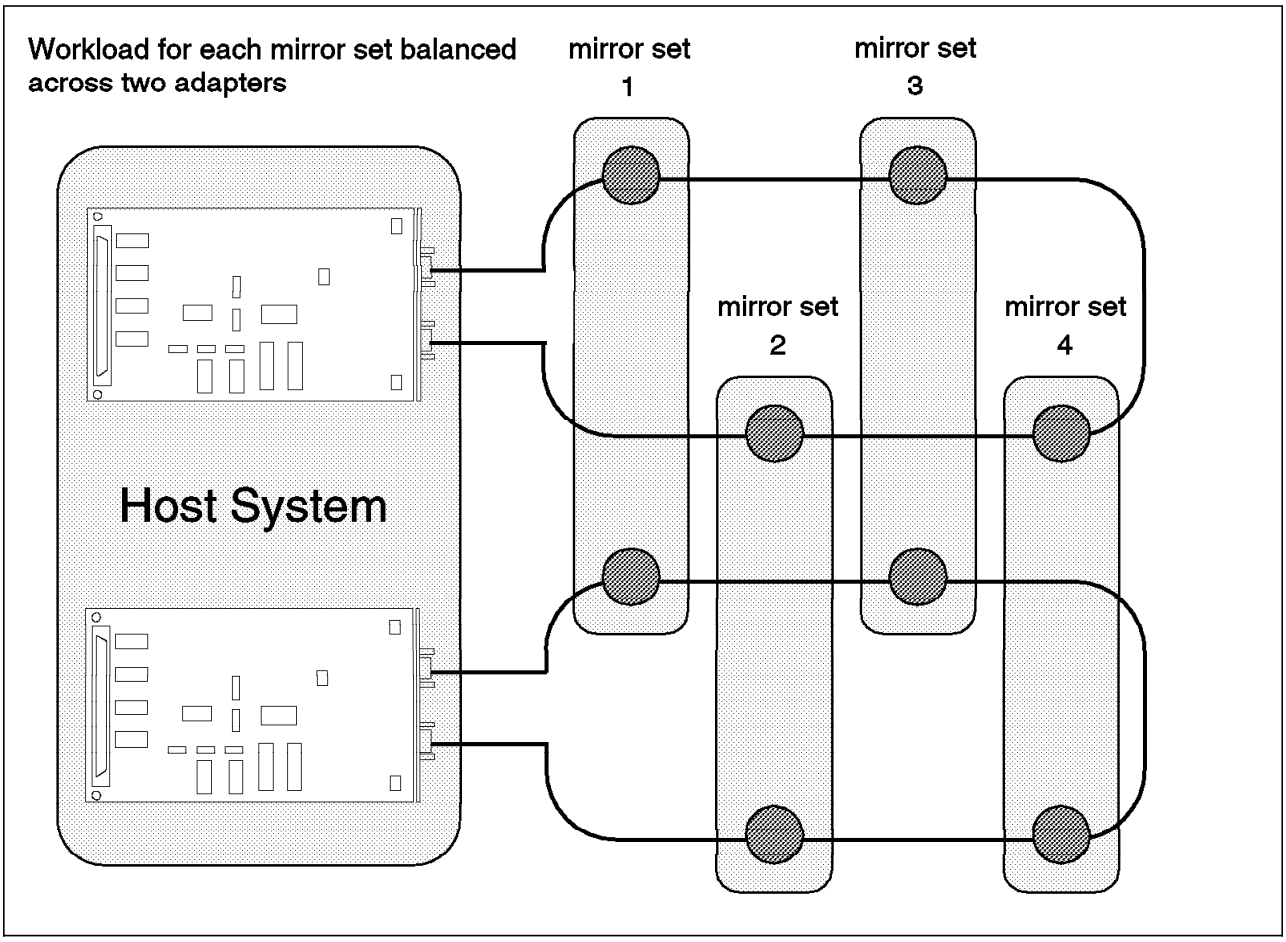

Figure 46. Mirroring across Adapters. This configuration is not vulnerable to a single adapter failure.

3.5.3 Planning for Growth

When planning your storage system, anticipate the need for future growth. Consider expansion within an existing enclosure or enclosures and how to add additional enclosures to your configuration.

When planning for growth within your storage drawers or towers, an issue that warrants careful consideration is how to arrange existing empty space in the enclosure. When using either the 7131-405 or the 7133, dummy disk drive modules must be placed in any slots that do not contain actual disk drives. The dummy disk drive modules provide circuit continuity to the SSA loop. The modules also maintain adequate cooling of the enclosure by preventing airflow from escaping through unoccupied disk drive slots. Do not place more than three dummy disk drive modules adjacent to one another.

When you use the 7131-405, you are required to place disk drive modules in slots 1 and 5. Slots 2, 3, and 4 can contain either a disk drive module or a dummy (see Figure 47 on page 54). Because only three slots are involved, there is really no right or wrong way to configure these slots.

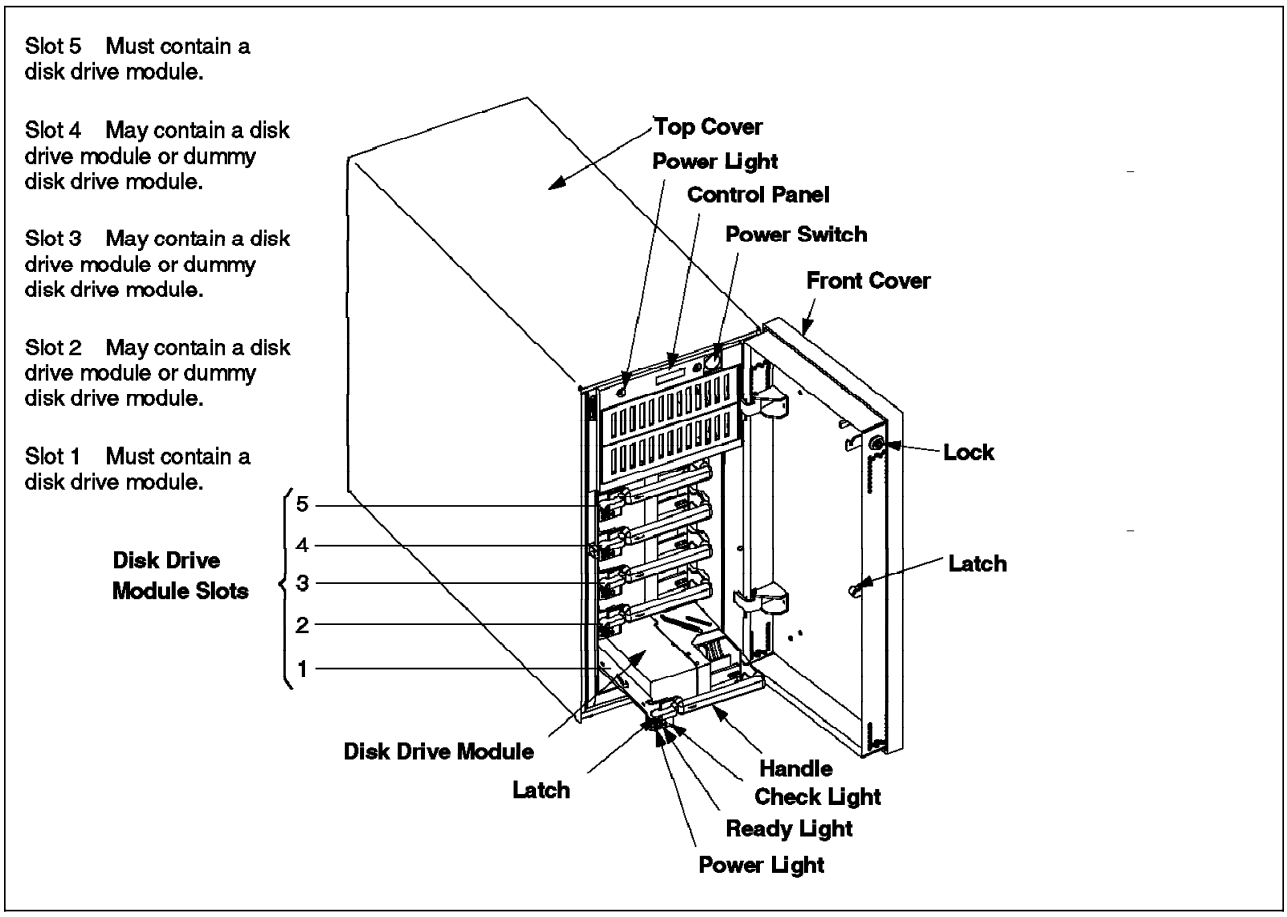

Figure 47. Arrangement of Dummy Disk Drive Modules in 7131-405.

When using dummy disk drive modules in the 7133, we recommend that you place the actual disk drives at the ends of each of the four strings within the enclosure. This means placing drives in slots 1, 4, 5, 8, 9, 12, 13, and 16 before placing drives in any other slots (see Figure 48 on page 55). The reason for this is simple: When cabling the subsystem, the attachment points correspond to these slots. When the subsystem is cabled as recommended, growth within the enclosure is straightforward: Simply replace the dummy disk drive modules with disk drives. Recabling will not be necessary. Figure 49 on page 56 demonstrates how this works.

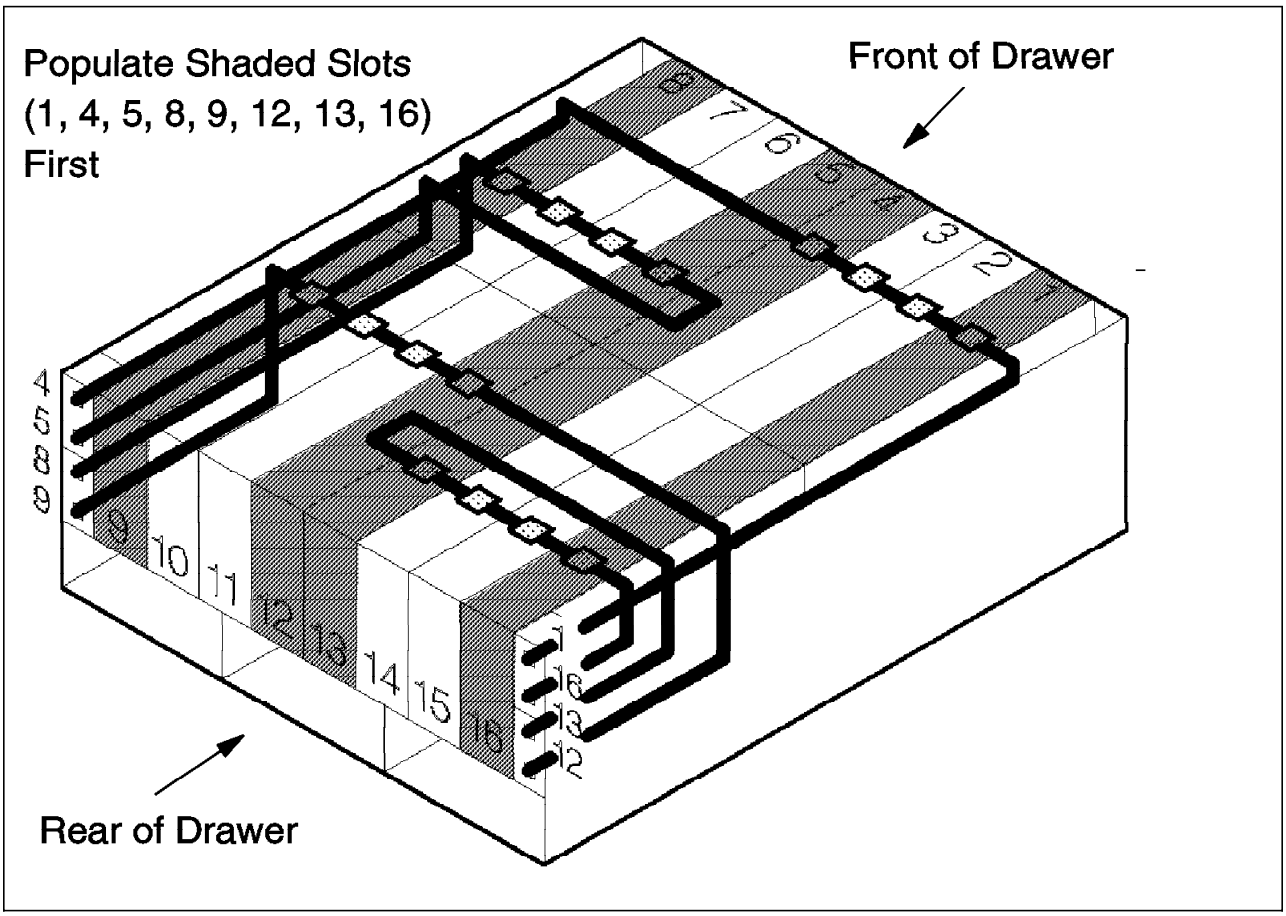

Figure 48. Preferred Order for Populating a 7133.

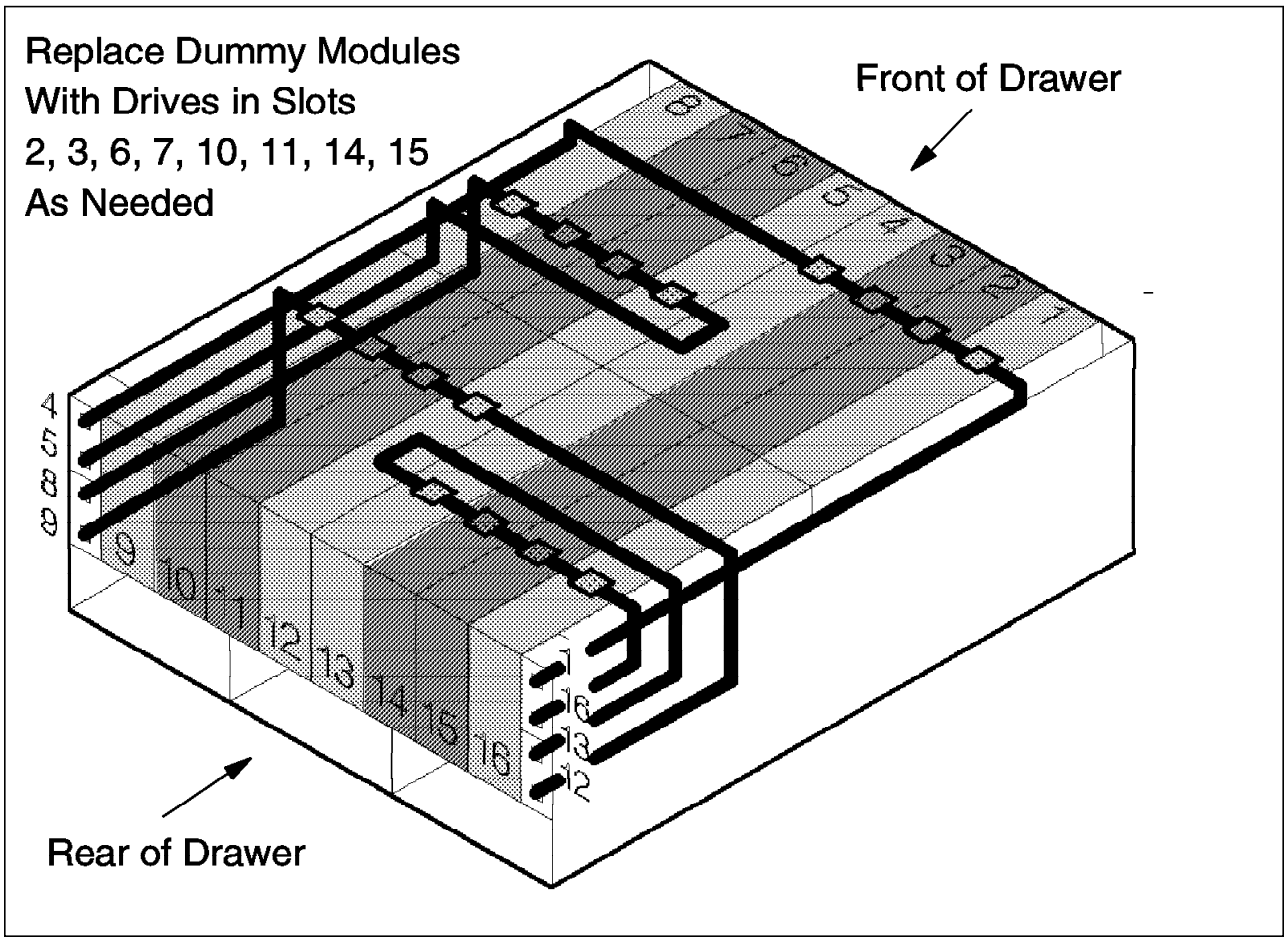

Figure 49. Replacing Dummy Disk Drive Modules in 7133. Recabling should not be required; just add the drives.

When arranging a storage system attached to a 7190, one characteristic that should be carefully examined is that the 7190 maps the SSA drives to SCSI ID/LUN combinations in the order encountered. Section 2.3.6, “SSA to SCSI ID/LUN Mapping” on page 21 discusses this in detail. The issue here is what happens to devices added after the 7190 has stored a configuration map in its nonvolatile memory. These new devices are assigned ID/LUN pairs, in the order encountered, starting with the next available ID/LUN pair beyond those already in the configuration map. For example, we will consider a situation where four new drives are added to an existing configuration of twenty drives (using LUN priority mapping). The existing drives are mapped to ID 0 LUN 0 through ID 2 LUN 4, respectively. The new drive nearest port 1 of the 7190 will be mapped to ID 2 LUN 5, the second new drive from port 1 will be mapped to ID 2 LUN 6, and so forth.

This characteristic can present a challenge if the 7190 is replaced. The new 7190 will consider all the drives to be new and map them in order of occurrence. This changes the ID/LUN pair assigned to any drives added after the initial configuration of the original 7190, and any drives that are farther from port 1 than the first of these added drives.

The workaround for this situation is relatively simple, although somewhat labor intensive. When replacing a 7190, temporarily remove any drives that were not in the original configuration of the original 7190. The new 7190 will map these drives to the same ID/LUN pairs as on the original 7190. Then, add the

additional drives to the configuration in the same order as they were added to the original configuration (before the 7190 was replaced). This will ensure that each drive retains the same ID/LUN pair.

3.5.4 Environment-Specific Considerations

Both the 7190 and SBus adapter offer the ability to leverage the power of one of the industry's leading storage technologies, SSA, in a multivendor environment. With this capability, you will encounter certain issues that are unique to a particular hardware or operating system environment.

3.5.4.1 UNIX Systems Compute Disk Size Differently

Disk capacities are often viewed in terms of the number of logical block assignments (LBAs) available on the disk. :12.logical block assignments Some UNIX implementations view two different disk drives of equal LBA capacities and different geometries (different numbers of heads and sectors per track) as being equal in size. Other implementations, including Solaris, calculate the capacity of disk drives in such a way that two equal-capacity drives of different geometries appear to have different capacities.

Why the Difference Occurs: The reason for the apparent difference in size arises from the way Solaris calculates capacities for disk partitioning. This calculation is done by reading the following items from the disk mode pages:

- Capacity (LBA)
- Number of heads
- Number of sectors per track

Solaris then calculates the number of cylinders on the disk, using the above information. It then uses the following formula to recalculate the capacity of the disk:

cylinders X heads X sectors per track

The difference occurs because Solaris uses an integer number of cylinders in the formula. In modern, advanced-technology disk drives, the actual number of cylinders is not always a whole number. Because of this, Solaris gets a different answer for different drive geometries even though the actual capacities of both drives are exactly the same number of bytes.

Disk Capacity Affects Replacement under RAID: When implementing SSA disk drives on Solaris in RAID and mirroring environments, you can encounter situations where drive replacement is difficult because of perceived differences in drive size. In RAID 1, the situation can occur where a full disk partition created on one 4.5 GB SSA drive cannot be mirrored on another 4.5 GB SSA drive, because the second drive appears to be smaller. A related situation occurs when you replace one 4.5 GB drive with another in a RAID 5 environment. The RAID software could reject the new disk because it appears to be too small, even though it actually has the same capacity. In RAID 0, replacing a 4.5 GB disk with an apparently smaller disk can reduce the effective size of the volume. If your storage design uses RAID 0 alone, this will have trivial impact, reducing your total array capacity by less than 5 MB. If your design uses RAID 0+1, the effect will be the same as with RAID 1, where you cannot establish or replace a mirror set member with an apparently smaller disk. For reference, Figure 50 on page 58 shows the differences in the calculated capacities of SSA disks.

Common Name	Starfire 1.1GB	Starfire 2.2GB	Starfire 4.5GB
Drive type	DFHC	DFHC	DFHC
Actual Capacity (LBA)	2,199,878	4,404,489	8,813,870
Actual Capacity (GB)	1.126338	2.255098	4.512701
Calculated #cyl	4071	4076	4078
Calculated #hd	4	8	16
Calculated #sec/trk	135	135	135
Calculated capacity (LBA)	2,198,340	4,402,080	8,808,480
Calculated Capacity (GB)	1.125550	2.253865	4.509942

Common Name	Scorfire 2.2GB	Scorfire 4.5GB	Scorpion 9.1GB
Drive type	DCHC	DCHC	DCHC
Actual Capacity (LBA)	4,404,489	8,813,870	17,796,077
Actual Capacity (GB)	2.255098	4.512701	9.111591
Calculated #cyl	2657	5320	5371
Calculated #hd	9	9	18
Calculated #sec/trk	184	184	184
Calculated capacity (LBA)	4,399,992	8,809,920	17,788,752
Calculated Capacity (GB)	2.252796	4.510679	9.107841

Figure 50. Actual and Calculated Capacities of SSA Drives.

Reducing Partition Size Avoids Problems: If your planned storage solution incorporates 4.5 GB drives in a RAID environment, there is a way to avoid the problems noted above. When configuring disk partitions for use by your RAID software, ensure that the entire amount of space occupied by all partitions is 5 MB less than the reported total capacity of the drive. Do this on each drive. Then, if you need to replace a drive, the worst case situation will be that Solaris thinks the new drive is smaller than the old drive but still of sufficient size to handle the same partitions. Drive substitution will be possible, regardless of how Solaris views the capacity.

For illustration, let's explore how you could use Veritas Volume Manager to implement this strategy. We assume that you want to place the entire physical disk under Volume Manager control, and we use encapsulation to accomplish this. Note that encapsulation requires two empty partitions, and two free cylinders at the beginning or end of the disk. We recommend that the free cylinders be placed at the end of the disk.

Using the Solaris format utility, create a partition that consists of the entire disk capacity, minus 5 MB, minus two additional cylinders. This gives you a 5 MB pad, plus the two free cylinders. Using a 4.5 GB Ultrastar 2XP drive, your partition would include cylinders 0 through approximately 5311, with cylinders 5312 through 5319 remaining. This remainder is 5.897 (rounded to 6) cylinders for the 5 MB, plus two cylinders. Now, use Volume Manager to encapsulate the disk. Volume Manager will create a two-cylinder private partition, plus your 5312-cylinder public partition which will appear as a Volume Manager virtual disk. Your 5 MB headroom remains untouched. Repeat this for all 4.5 GB

drives, and use the large partition to create your RAID 1 mirror sets, RAID 5 arrays, or RAID 0 arrays.

3.5.4.2 Where to Place the 7190 Service Functions Partition

The 7190 service functions application provides operations and maintenance capabilities such as downloading disk drive microcode, displaying disk drive vital product data (VPD), performing diagnostics on the 7190 or disk drives, and so forth.

The application requires at least one partition of at least one disk. This partition stores error logs for use by the service functions application. The partition must be at least 1 MB in size and must be available for the exclusive use of the 7190—any other data placed in the partition will be overwritten. To maintain availability of the error logs in the event of a disk failure, we recommend that you create this dedicated partition on more than one disk. After doing so, the dedicated partitions must also be referenced in the 7190.cfg file, which allows the service functions application to locate the partitions. The 7190.cfg file is to be placed in the same directory as the service functions application so that the application can locate it. A sample 7190.cfg is shown in Figure 51.

```
# IBM 7190 service function configuration file
# please assign the service function's partition
# example :
# Name Partition
#nnn    /dev/rdsk/c1t0d0s7
#nnn    /dev/rdsk/c1t0d1s7
#
SSA1    /dev/rdsk/c5t0d0s7
SSA1    /dev/rdsk/c5t0d1s7
SSA2    /dev/rdsk/c6t0d0s7
SSA2    /dev/rdsk/c6t0d1s7
```

Figure 51. 7190.cfg Example. This shows a 7190.cfg for two 7190s, SSA1 and SSA2, attached to Controllers 5 and 6, respectively. Each references a service function partition on Slice 7 of ID 0 LUNs 0 and 1, attached to its respective controller number.

When deciding where to create these dedicated partitions, you must also consider the impact of logical volume management software like Solstice DiskSuite or Veritas Volume Manager. Under Volume Manager, in particular, three separate options exist. The first option would be to create the partitions on disks that are not under Volume Manager control.

If it is not feasible to use disks beyond the control of Volume Manager, a second option would be to use Volume Manager to encapsulate the disk or disks involved. After doing so, use the raw partition of the VM disk or build a simple plex (a VM name for mirror set member) and volume on the VM disk. Use the raw logical volume device name in the 7190.cfg file. An example of this device name is /dev/vx/rdsk/groupname/volname.

A third option would be to create Volume Manager subdisks for the respective partitions. A subdisk could then be used to create a simple plex. The plex, in turn, can be used to build a volume that can be referenced by the 7190.cfg file. We recommend that you exercise caution when using this option, as the Volume Manager Basic Operations menus operate only with entire physical disks. The first two options of providing a private partition would be preferable.

3.5.4.3 Bypass Card Configuration

A special topic for consideration is how to configure a 7133-020 or 7133-600 enclosure that is connected to more than one host system. These 7133 models incorporate bypass circuits, which are designed to preserve the integrity of the SSA loop by allowing the bypass of an unavailable host, adapter, or cable in the loop. The circuits are present on each adjacent pair of SSA ports on the enclosure and connect the ports to the respective drive slots inside the enclosure (see Figure 52). The circuits operate in two modes: *bypass mode* (the default setting) and *forced inline mode*.

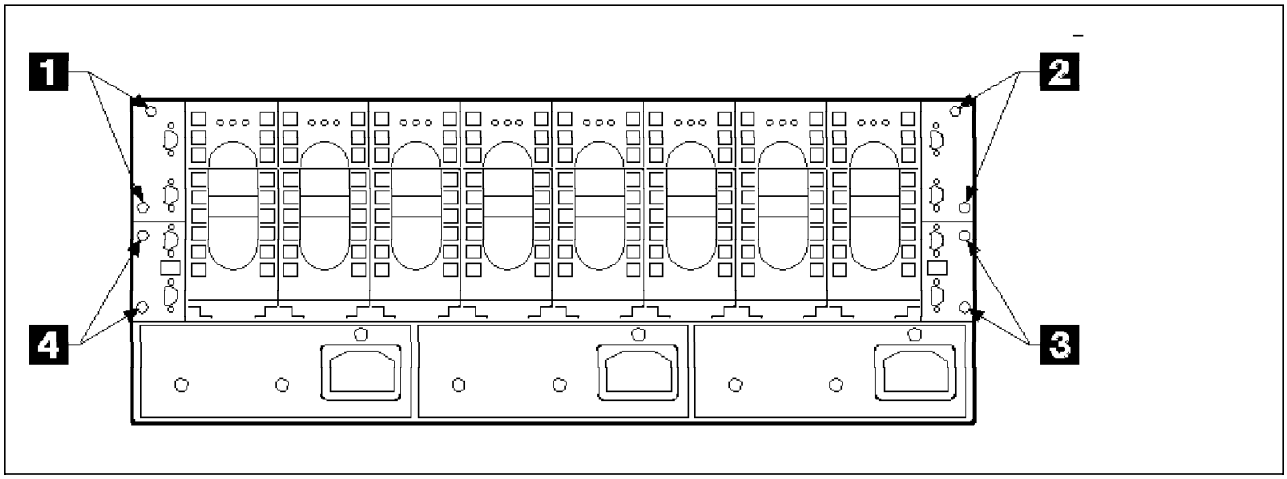

Figure 52. Location of Bypass Circuits in 7133.

The function of a bypass circuit is quite simple: in bypass mode, the circuit monitors both of its external ports. The circuit remains open (*inline* state) as long as power is detected at one of the two ports. Power at either port means that a cable is attached with a powered adapter or enclosure at the other end. In an inline state, the ports remain connected to the associated drive slots inside the enclosure (see Figure 53 on page 61).

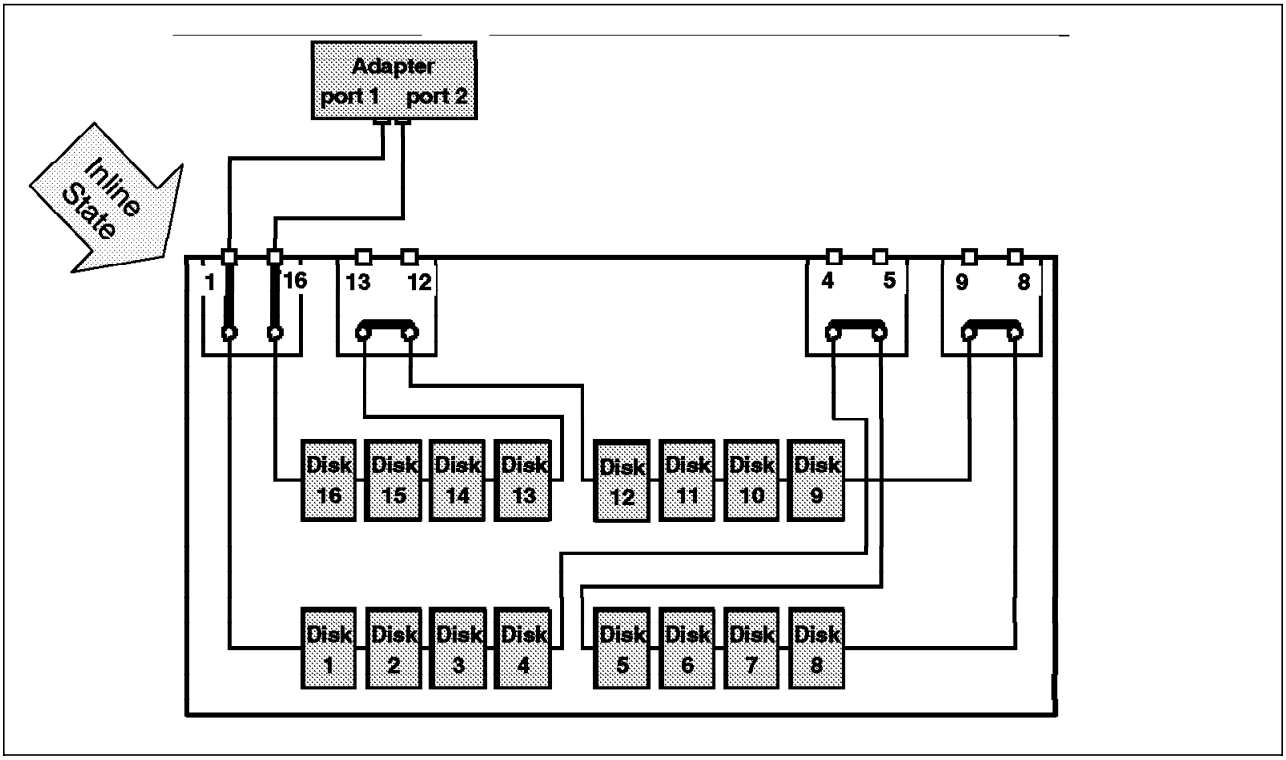

Figure 53. Inline State. Ports are connected to associated drive slots.

If power is not sensed at either port, it means that either both cables have been removed or both cables are connected to an adapter, adapters, or other enclosures that are powered off. In this situation the circuit closes, disconnecting the external ports and connecting the two ports internally as if the drive slots were cabled to each other, preserving the continuity of the loop. This is known as *bypass* state and is illustrated in Figure 54 on page 62.

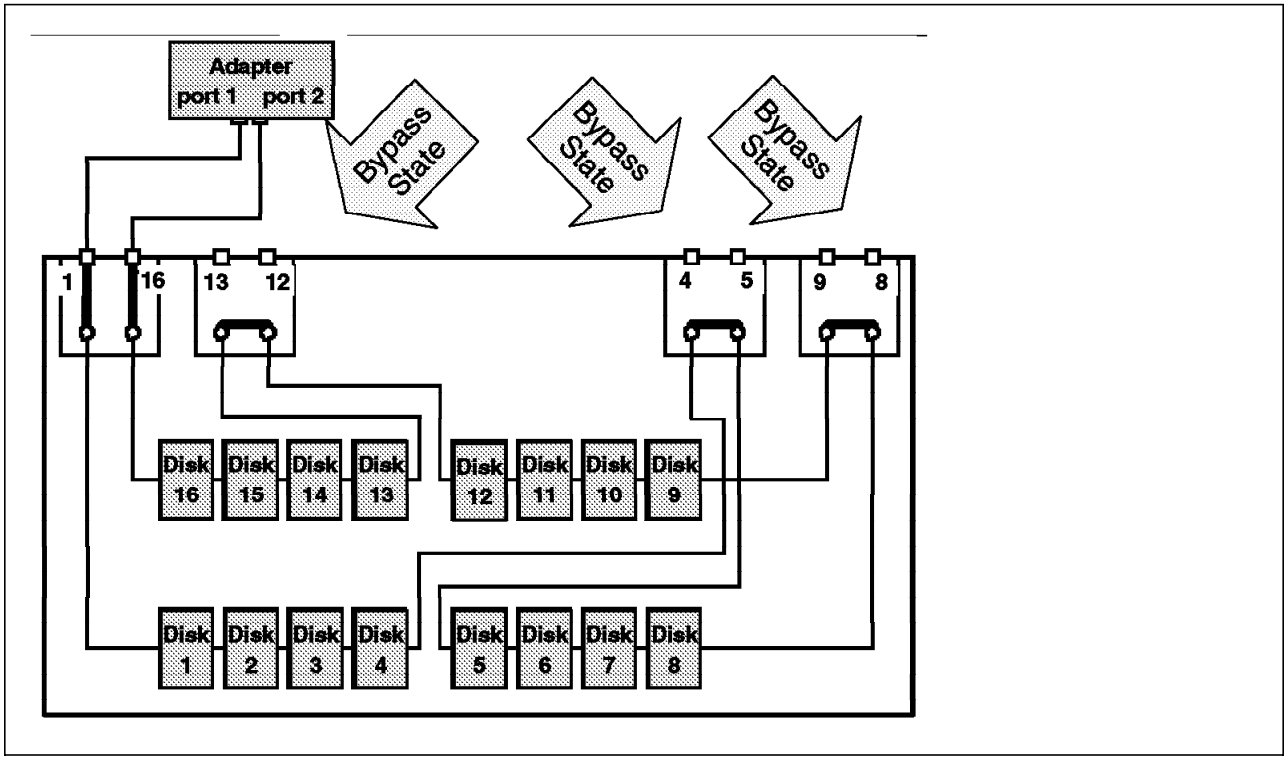

Figure 54. Bypass State. Adjacent drive slots are disconnected from their external ports and connected to each other.

When the bypass circuit is placed in forced inline mode, the switching ability of the circuit is disabled. In this mode, the circuit is always in the inline state: the external ports are always connected to the respective drive slots, regardless of whether power is detected at either port.

The important issue is how to utilize the bypass circuits to your advantage. Although designed for a multihost multiadapter configuration, the circuits can cause problems if not cabled properly. To evaluate the configuration of each bypass circuit in your storage subsystem, we offer the following rule-of-thumb: If closing a particular bypass circuit (placing it in bypass state) results in a valid configuration, enable that circuit. If closing a particular bypass circuit causes either an unknown configuration or one that is known to be invalid, place that circuit in forced inline mode. See the *Installation and User's Guide* for your respective enclosure for information about configuring the enclosure for forced inline mode or bypass mode.

3.5.4.4 7190 and SBus Adapter Addressability

As discussed in Chapter 2, "SBus Adapter and 7190 Technology" on page 9, SSA drives attached to Solaris systems through the 7190 and SBus adapter appear to the host as path names using SCSI controller, target, LUN combinations. If your storage plans include using multiple adapters to connect a loop to a host, some special consideration is necessary. This configuration will result in more than one path name pointing to the same drive, making it appear as two different drives. Care must be taken to avoid trying to access the drive through both path names at the same time. Ensure that you have appropriate software in place to properly control this environment.

Another addressability situation, unique to the 7190, is that the adapter itself appears to the system as one or more SCSI targets, depending upon the mode

selected and the number of SSA drives attached. See 2.3.6, “SSA to SCSI ID/LUN Mapping” on page 21 for a detailed explanation of SCSI ID/LUN mapping. In short, the adapter will present the SCSI controller with as many ID/LUN combinations as there are SSA drives, starting with ID 0 LUN 0. In LUN priority mapping, this will consume one target ID (starting with zero) for every eight disks attached. In ID priority mapping, one target ID will be used for each disk attached.

While not recommended for performance reasons (see 3.5.1.3, “Place 7190 Alone on SCSI BUS” on page 43), you may wish to install the 7190 on a SCSI bus with other devices. If you choose to do this, note that no other devices on the bus can use the IDs occupied by the 7190 because each target ID on the bus must be unique. If two devices with the same target ID occupy the same SCSI bus, unpredictable (and probably, undesirable) results will occur. Since the 7190 will start mapping from ID 0 LUN 0, you must configure the 7190 with all disk drives before adding any additional SCSI devices to the bus. Then, when adding devices, you must avoid the target IDs used by the 7190 (and, of course, the SCSI controller itself). It becomes evident that sharing a SCSI bus between a 7190 and other devices can be quite complicated, which is another good reason for not doing so.

Chapter 4. Configuration and Installation

The procedures and guidelines in this chapter will help you prepare for installation and configuration of the IBM SSA Interface Controller for Sun SBus and the 7190. This chapter includes the following topics:

- IBM SSA Interface Controller for Sun SBus
- IBM 7190 model 100

4.1 IBM SSA Interface Controller for Sun SBus

The IBM SSA Interface Controller for Sun SBus enables you to attach IBM 7133 and 7131 storage subsystems to specific Sun host systems. The adapter is configured as a standard SBus card consisting of a printed circuit board with two SSA ports and SBus SSA logic. Communication between the host system and the adapter is through the SBus interface. Figure 55 shows an IBM SSA Interface Controller for Sun SBus.

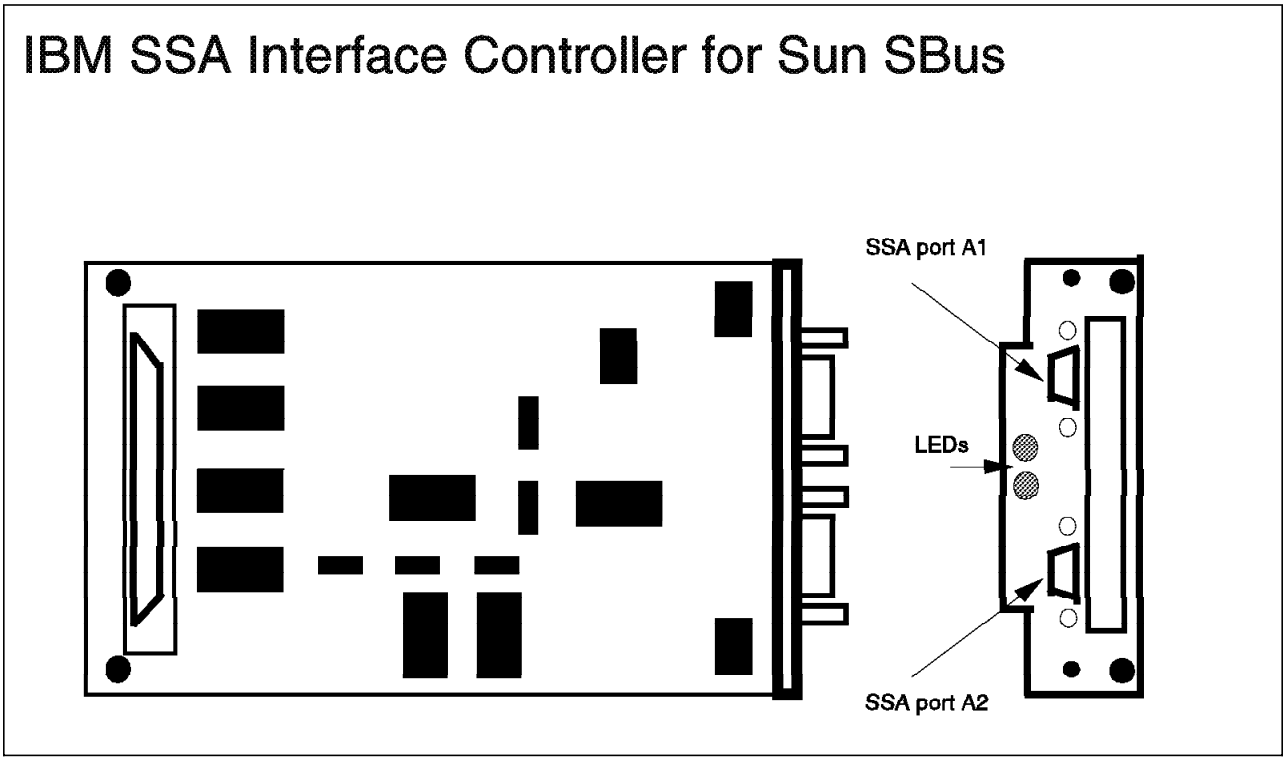

Figure 55. IBM SSA Interface Controller for Sun SBus.

In this section we discuss how to prepare for installation and describe the installation procedure itself.

4.1.1 Preparing for Installation.

Before you start the installation, ensure that you have the following tools, information, documentation Web sites, and hardware and software:

Tools

You may need a small, flat-head screwdriver to secure the SSA connectors to the adapter. It is possible to use the special screwdriver supplied with the 713x to tighten the SSA connector screws.

Information

Obtain the following information from your system administrator:

- Host system and slot available for the adapter card
- Host system software requirements. Review these requirements in Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27.
- IBM 7131 or 7133 Disk Subsystem Installation plan. It may be necessary to have an IBM service representative install the disk subsystem.
- SSA loop configuration. Refer to Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27 for information about loop configurations and cluster design.
- Installation/test verification plan. These are tasks to perform when the installation is completed, for example disk configuration, file system definitions, database installation, database definition, load data.

Documentation and Web Sites Obtain the following from system hardware documentation:

- SBus card installation procedure
- *IBM SSA Interface Controller for Sun SBus: Installation and User's Guide*
- Sun host system installation manual You may also want to look at these reference manuals:
  - *7133 SSA Disk Subsystem for Open Attachment: Installation and User's Guide , SA33-3273*
  - *7133 SSA Disk Subsystem for Open Attachment: Service Guide, SY33-0191.*
  - *Translated Safety Notices for Open Attachment: GC26-7246.*
  - *7131 Model 405, SSA Multistorage Tower for Open Attachment: Service Guide, SY32-0405.*
  - *7131 Model 405, SSA Multistorage Tower for Open Attachment: Installation and User's Guide, SA26-2001.*

  Additional information is available on the following Web sites:

  - IBM Storage
    Visit the SUPPORT pages at the IBM Storage Home Page:
    http://www.storage.ibm.com
  - Sun Microsystems
    SunSolve Online Free Services Web Page:
    http://sunsolve.sun.com/sunsolve/contractservices.html
  - Veritas Technology Solutions Ltd.
    Visit the Veritas Home Page
    http://www.veritas.com

  Hardware and software

  Verify that the following are in perfect condition:

- IBM SSA Interface Controller for Sun SBus adapter card
- CD-ROM with device driver

If you observe shipping damage, do not install the adapter without IBM approval. Report all observed damage to your sales representative.

4.1.2 Installation Procedure

In this section, we explain the procedure to install the IBM SSA Interface Controller for Sun SBus. We cover the following topics:

1. Safety information
2. Adapter card installation
3. Software installation
4. Assigning mapping address
5. Utilities description
6. Drive configuration
7. Configuring multi-initiator clusters

4.1.2.1 Safety Information

Lethal voltages may be present in your system even with the power switch off. The following safety notices are used in this installation procedure. Translations of these notices are in *Translated Safety Notices for Open Attachment, GC26-7246*, included with this product.

SAFETY NOTICES

- Please follow all electrical safety instructions described in your system documentation.
- An electrical outlet that is not correctly wired could place hazardous voltage on metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock.
- Before installing or removing parts, ensure that the power cables for the system unit and all attached devices are unplugged. Electrostatic discharge (ESD) protection is required for your system when its cover is removed, and for the adapter or other hardware components at all times.

4.1.2.2 Adapter Card Installation

In this section we describe the procedure for installing the adapter card in your host. This procedure assumes that your host system is properly installed and operating according to the host system manuals. You must ensure that the system hardware and system software meet the requirements defined in Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27..

Before you begin, inspect the shipping container. Refer to the document included with the order and verify that you have all of the parts indicated. If the shipping container or its contents are damaged or parts are missing, contact your IBM sales representative.

We recommend that you back up or copy all files before making major changes to the system.

If the system is on and operating, perform the shutdown procedure and turn the host system switch to off.

You are ready to install the adapter card. Follow these steps.

Installing the IBM SSA Interface Controller for Sun SBus Adapter Card: See Figure 56

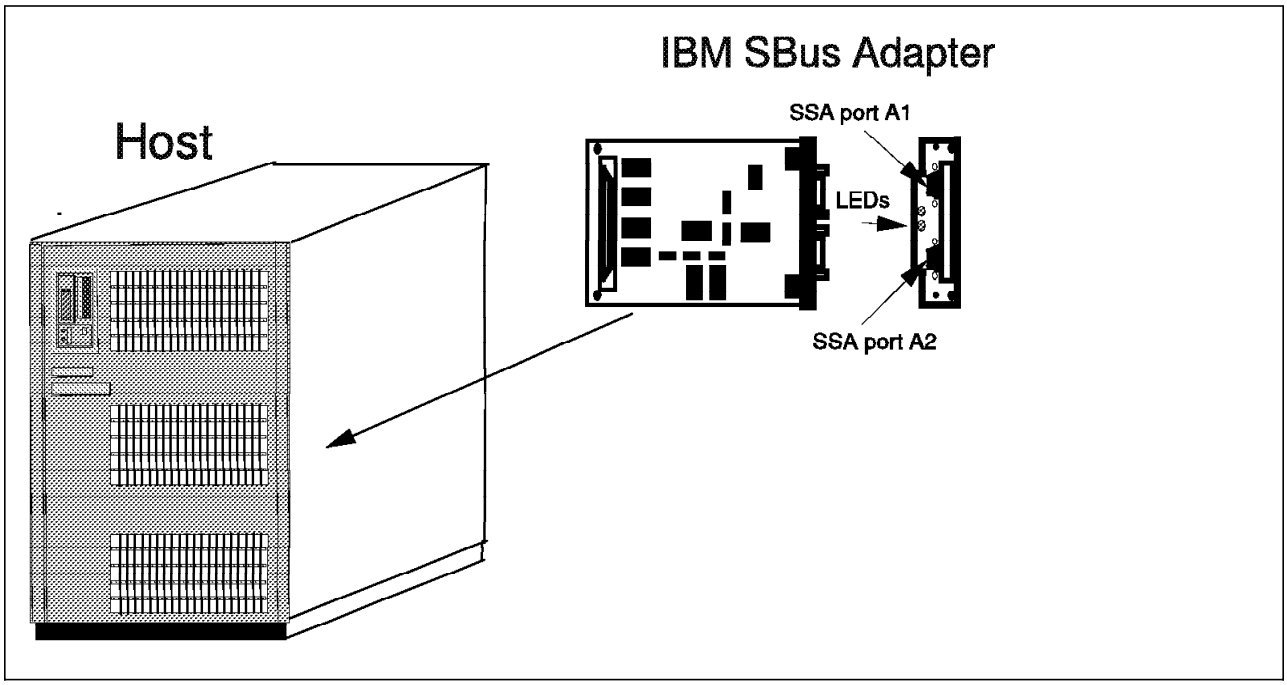

Figure 56. Installing IBM SBus Adapter Card

The Sun host system installation manual contains instructions on how to install the SBus card in the host system board. Install the adapter card, following the instructions in the manual. Make sure that the adapter is firmly seated in the SBus slot. Replace the system cover.

Now you are ready to install the 713x SSA disk storage subsystem.

Installing the 713x SSA Disk Storage Subsystem: An IBM service representative may be required to install the 713x SSA disk storage subsystem.

Follow these steps:

1. Turn off the 713x SSA disk storage subsystem unit.
2. Check the power cables and make sure that they are removed from the power outlet.
3. Install the SSA cables between the adapter card and the SSA disk storage subsystem as shown in Figure 57 on page 69

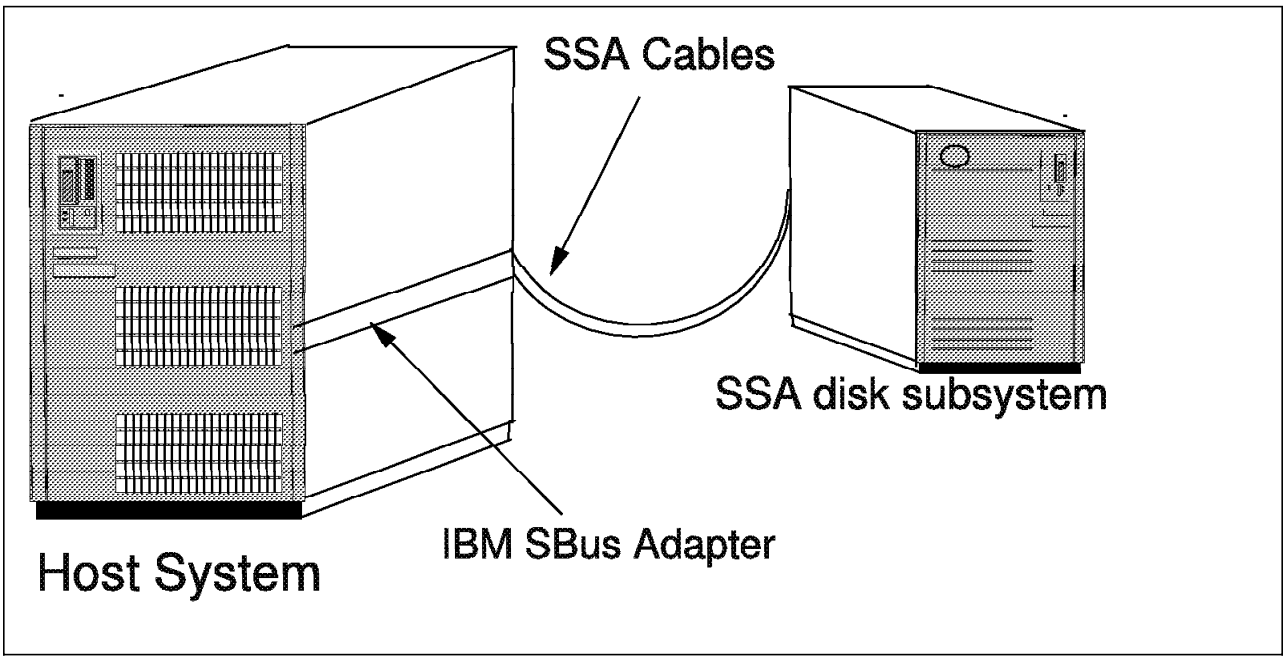

Figure 57. Installing 713x Disk Storage Subsystem

To make an SSA loop, you need at least two external SSA cables and at least one SSA device. Connect one end of one cable to either port on the adapter and the other end to one of the ports on the device. Then connect one end of the other cable to the second port on the adapter and the other cable end to the other port on the device. The Figure 58 shows a loop connection.

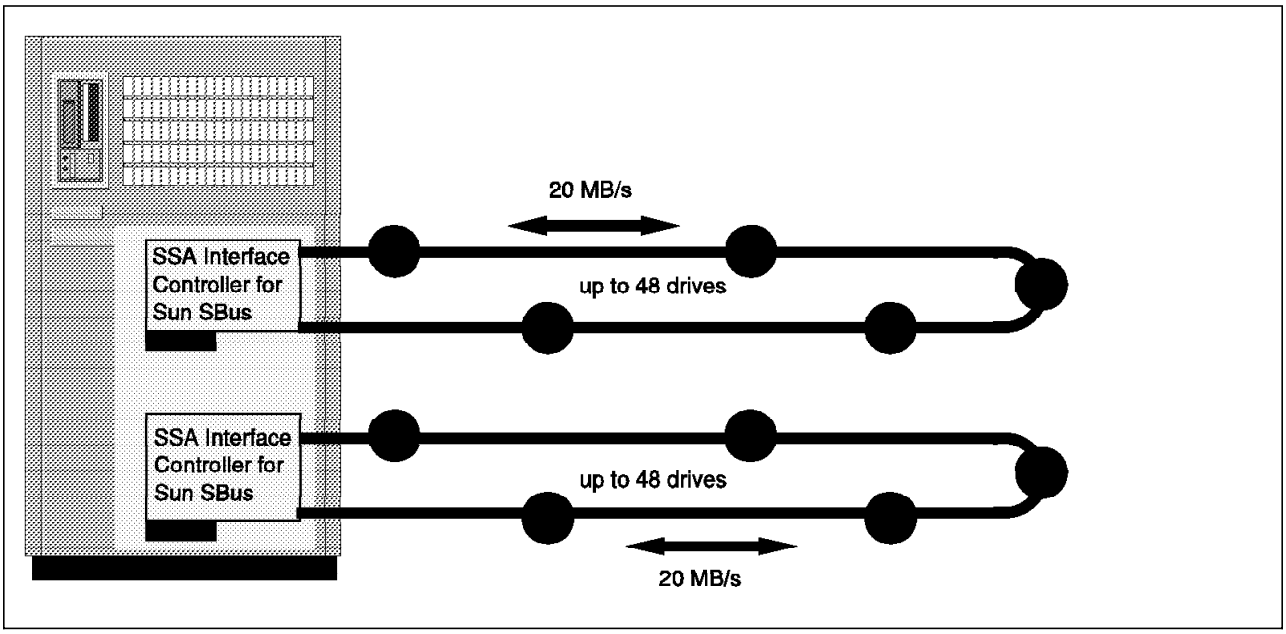

Figure 58. SBus controllers in loop topology

Multiple devices can be added to the loop. For each device added, you need additional external SSA cables.

After you have set up the devices and cables in the configuration you want (see Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27). Make sure the cables are tightly screwed in to the device and adapter ports.

Now are you ready to power up the SSA devices and your system.

Here is the power-up sequence that we recommend:

1. Power up the 713x SSA storage disk subsystem. Set the 713x storage disk subsystem power switch to on.
2. Power up the system.

Use the following procedure to verify the status of the devices:

1. Check the LEDs on all drives in the 713x SSA storage disk subsystem:
   - Both green LEDs on all drives should be on (steady, not blinking).

   If the LEDs are not in the state indicate above, refer to the following documentation to identify the condition:
   - *7133 SSA Disk Subsystem for Open Attachment: Service Guide, SY33-0191.*
   - *7131 Model 405, SSA Multistorage Tower for Open Attachment: Service Guide, SY32-0405.*

   If assistance is required, call your IBM service representative.
2. Check the LEDs on the SBus adapter. After power-on, the adapter requires 2 s to complete the power-on self test (POST). When the POST is complete, the green LED should be on and steady, not blinking.

4.1.2.3 Software Installation

In this section, we describe the process of copying the files to disk and installing the packet driver in Solaris. The driver has two major components: the SSA-compliant host bus adapter, which communicates with the operating system, and the Serial Driver Interface (SDI), which communicates with the SBus adapter. Both components also communicate with each other.

The host bus adapter driver complies with the level of kernels specified by Sun Microsystems, Inc.: Sun Common SCSI Architecture (SCSA), Device Driver Interface (DDI) and Driver Kernel Interface (DKI).

SCSA defines the interface between a SCSI target driver and a host adapter driver. DDI and DKI (usually written as DDI/DKI) define the kernel services available to a device driver.

The proper use of these interfaces enables a driver to correctly operate in various Sun systems, regardless of whether the particular system contains one or more processor modules. The host adapter driver is therefore SMP-capable. These interfaces are only usable at the driver level.

The procedure for installing the driver package requires root privileges. Follow these steps to install the driver package:

1. Load the CD-ROM into the host system. Wait 10 seconds to allow the system to perform the physical mount.
2. To install the driver package we use the pkgadd Sun utility to allow software installation. Enter this command:

   ```
   pkgadd -d  /cdrom/unnamed_cdrom/SSASBUS.PKG
   ```

where /cdrom/unnamed_cdrom is the mount point for the CD-ROM, and SSASBUS.PKG is the name of the package you are installing as indicated on the CD-ROM label (Figure 59 on page 71 shows a sample installation).

```
root>pkgadd -d  /cdrom/unnamed_cdrom/sl_1_30.pkg

The following packages are available:

1 PTLsunl     sunlight device driver and utilities

                     (Solaris2.4,sparc) 1.30


Select packag(s) you wish to process (or 'all' to process

all packages). (default: all) (?,??): >cr<

Processing package instance <PLTsunl> from </cdrom/sl_1_30.pkg>

sunlight device driver and utilities (Solaris2.4,sparc) 1.30

Copyright 1997, Pathlight Technology, Inc.

All rights reserved.

## Processing package information

## Processing system information.

## Verifying disk space requirements.

## Checking for conflicts with packages already installed.

## Checking for setuid/setgid programs.


This package contains scripts which will be executed with
super-user

permission during the process of installing this package.


Do you want to continue with the installation of <PTLsunl>
(y,n,?) y


Installing sunlight device driver and utilities as <PTLsunl>


## Installing part 1 of 1.
```

Figure 59 (Part 1 of 2). Sample Installation

```
/kernel/drv/sunlight

/kernel/drv/sunlight.conf

/opt/PTLsunl/firmware/a74sc081.bin

/opt/PTLsunl/firmware/fw280.aout

/opt/PTLsunl/firmware/irrst084.bin

/usr/include/sys/scsi/adapters/shstio.h

/usr/bin/ssau

/usr/bin/sscf

"Verifying class <none> "

Modifying /etc/devlink.tab

Modifying /kernel/drv/sd.conf

"Verifying class <sed> "

## Executing postinstall script.


Installation of <PTLsunl> was successful..
```

Figure 59 (Part 2 of 2). Sample Installation

Figure 59 on page 71 shows a typical output from the pkgadd command.

When you enter the pkgadd command, the result is to verify the CD-ROM unit and extract the package information from the device. A prompt with the packages that can be installed into the system then appears, as follows:

```
The following packages are available:

1 PTLsunl     sunlight device driver and utilities

                     (Solaris2.4,sparc) 1.30


Select packag(s) you wish to process (or 'all' to process

all packages). (default: all) (?,??): >cr<
```

The system waits for the user response. Press Enter to accept the default.

The system then validates the package information, checks dependences with other system software, checks the package security levels (user level and group level), and displays the next screen:

```
sunlight device driver and utilities (Solaris2.4,sparc) 1.30

Copyright 1997, Pathlight Technology, Inc.

All rights reserved.


## Processing package information

## Processing system information.

## Verifying disk space requirements.

## Checking for conflicts with packages already installed.

## Checking for setuid/setgid programs.
```

This software has user level and group level security system. Only the superuser (root) can perform the installation process. The system shows the following information and waits for user response:

```
This package contains scripts which will be executed with
super-user permission during the process of installing this
package.

Do you want to continue with the installation of <PTLsunl>
(y,n,?) y
```

Respond **y** to continue with the installation

The system begin the installation and shows the process result in Figure 60 on page 74.

```
Installing sunlight device driver and utilities as <PTLsunl>

## Installing part 1 of 1.

/kernel/drv/sunlight

/kernel/drv/sunlight.conf

/opt/PTLsunl/firmware/a74sc081.bin

/opt/PTLsunl/firmware/fw280.aout

/opt/PTLsunl/firmware/irrst084.bin

/usr/include/sys/scsi/adapters/shstio.h

/usr/bin/ssau

/usr/bin/sscf

"Verifying class <none> "

Modifying /etc/devlink.tab

Modifying /kernel/drv/sd.conf

"Verifying class <sed> "

## Executing postinstall script.
```

Figure 60. Installation Process for PTLsun1

When the process finishes, the system shows the successful result:

Installation of <PTLsunl> was successful.

If the File Manager is running, it opens a window that displays the mount point and the contents of the CD-ROM. If the File Manager is not running, you can use the df command to display the mount point.

3. If the File Manager is running, select **Eject Disk** to unmount and remove the CD-ROM. If the File Manager is not running, at the command line enter the **eject cdrom** command to unmount and remove the CD-ROM.
4. Reboot from the Solaris prompt with the **reboot -- -r** command. This command halts the operating system, synchronizes the superblock, and restart the machine. This process allows the new driver to be recognized as part of the operating system and the devices attached to SBus controller can be configured.

   The superblock is a global structure that contains information about the file system, describes how many blocks need to be kept free, the number of disk revolutions per second, the maximum number of blocks a file can use in a cylinder group, total number of cylinders per cylinders group, and the number of inodes per cylinder group.

An inode is a background structure that directs the operating system to the location of the data blocks that are stored on the disk fo a single file. The inode contains pointers (address) that tell the operating system how to find the data blocks for its file.

Use of reboot -- -r command avoids file system corruption.

4.1.2.4 Assigning Mapping Addresses

In this section we describe the process of assigning a mapping address to each of the SSA devices in the adapters. Mapping is needed because of the use of the SCSI address instead of the SSA unique identification number.

The host bus adapter device driver for the IBM SSA Interface Controller for Sun SBus adapter is designed to work with the Solaris operating system. The device driver supports up to two IBM SSA Interface Controllers for Sun SBus adapters installed in the same system. Each adapter contains two SSA port connectors on its panel. Each adapter supports up to 48 SSA devices in either a loop using both ports or in one or two strings (one string is possible in each port). The two adapters in a system must not be connected to the same SSA web. There can be four initiators within a single SSA web. Figure 61 shows IBM SSA Interface Controller for Sun SBus cards in a loop topology.

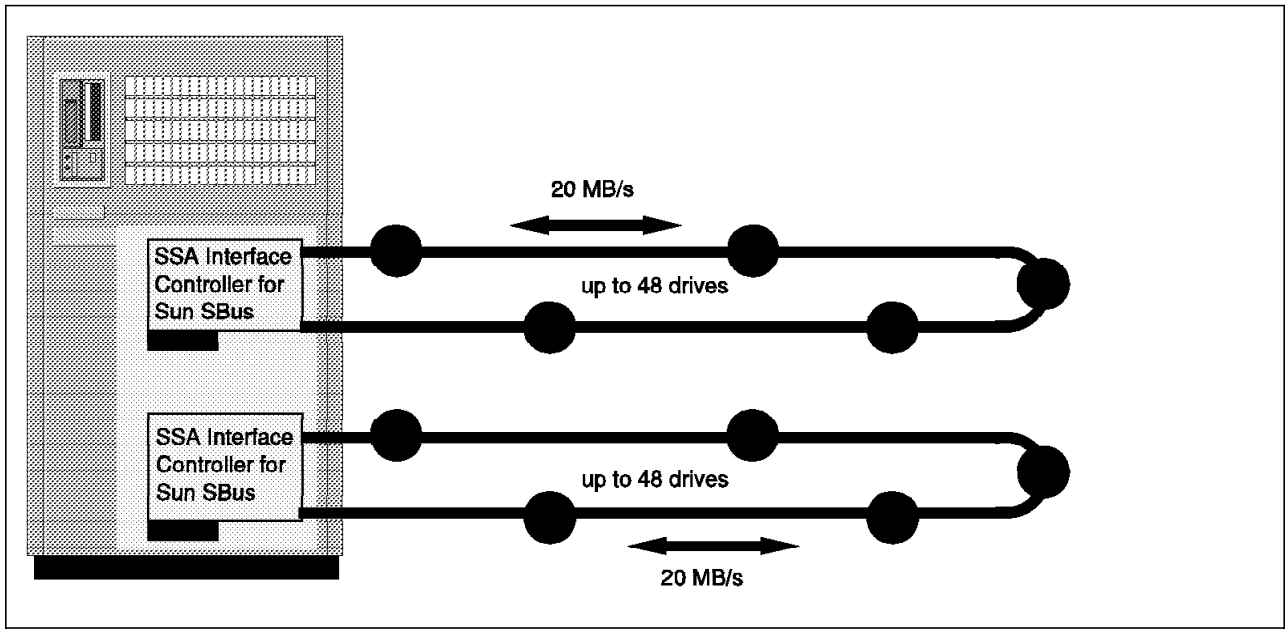

Figure 61.

For additional information about configurations, limitations and concepts refer to Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27.

After the driver files have been installed on the disk, if any IBM SSA Interface Controllers for Sun SBus are detected on the SBus, the driver is automatically installed in the operating system when the system boots.

The process of automatic installation or “autoconfiguration” consists of:

1. Loading the driver code from disk into memory
2. Registering the driver with the system

3. Identifying the SSA devices on the adapter's web

The driver appears to the operating system as a driver for a SCSI adapter. Thus both the operating system and user applications can communicate with SSA devices as if they were SCSI devices.

The SCSI target driver, **sd**, which is a standard component of Solaris, provides SCSI function.The SCSI target driver interprets all read and write requests targeted to the SSA disk and then passes the request to the adapter, which in turn forwards it to the specified SSA disk.

The unique identification number (UID) identifies an SSA device and a SCSI ID identifies a SCSI device. The operating system accepts the SSA disks as SCSI disks because a mechanism maps from SCSI IDs to SSA UIDs.

Each time the system is booted, the Solaris operating system probes for all attached devices. In the case of SSA devices that masquerade as SCSI devices, the system probes for all devices on each installed SSA adapter. The probe can proceed only if each installed SSA device is assigned a SCSI ID.

The host adapter creates a default configuration map by using an internal algorithm to assign default SCSI IDs to drives. The adapter firmware finds each drive on the SSA web.

The algorithm assigns SCSI IDs based on the order in which each drive is found on the SSA web by the adapter firmware. For example, the first device found is assigned SCSI ID 00, the second device ID 01, and so on. Figure 62 shows a logical web map.

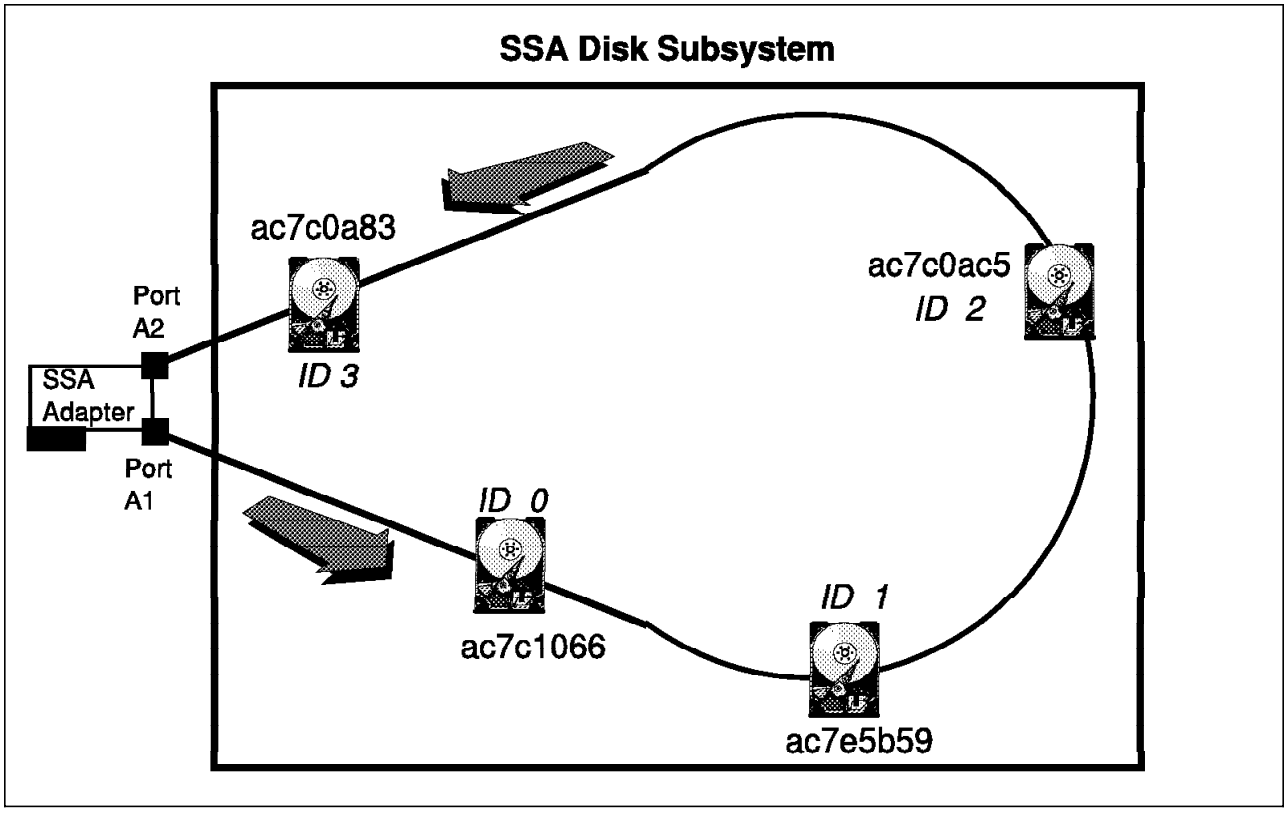

Figure 62. Logical Web Map

Figure 62 shows a default autoconfiguration map.

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID        SCSI

#High         Low        ID

0800          ac7c1066   00

0800          ac7c5b59   01

0800          ac7c0ac5   02

0800          ac7c0a83   03
```

Figure 63. Default Autoconfiguration Map

The order in which devices are found is not necessarily the same as the order in which they are cabled together. This is especially true if the devices are configured as two strings during one boot and a single loop during the other boot. In this case, if the host adapter driver does not have other instructions, a new SCSI ID may be assigned to the devices at each boot. This method of persistent configuration ensures that the SCSI ID is persistent across boots. Figure 64 on page 78 shows a logical web map loop and Figure 65 on page 78 shows a logical web map string.

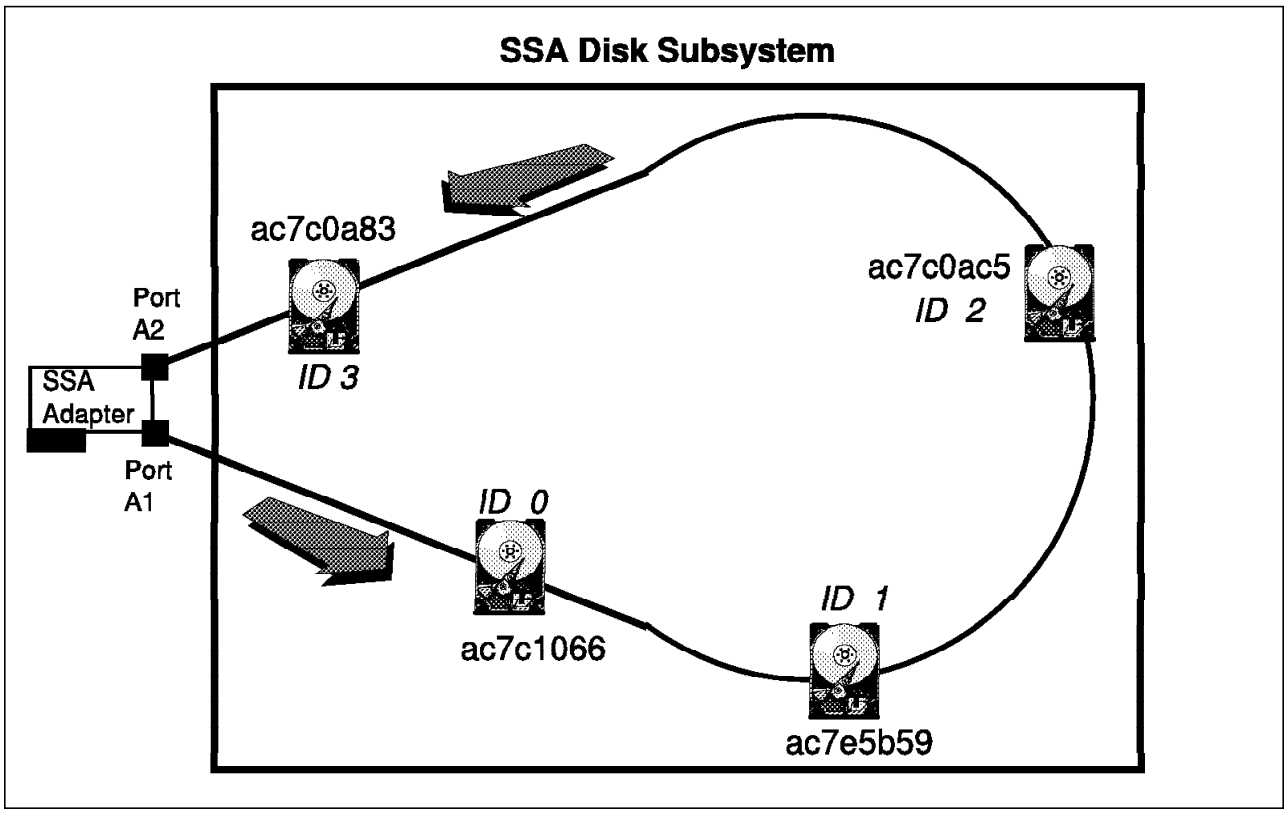

Figure 64. Logical Web Map - Loop

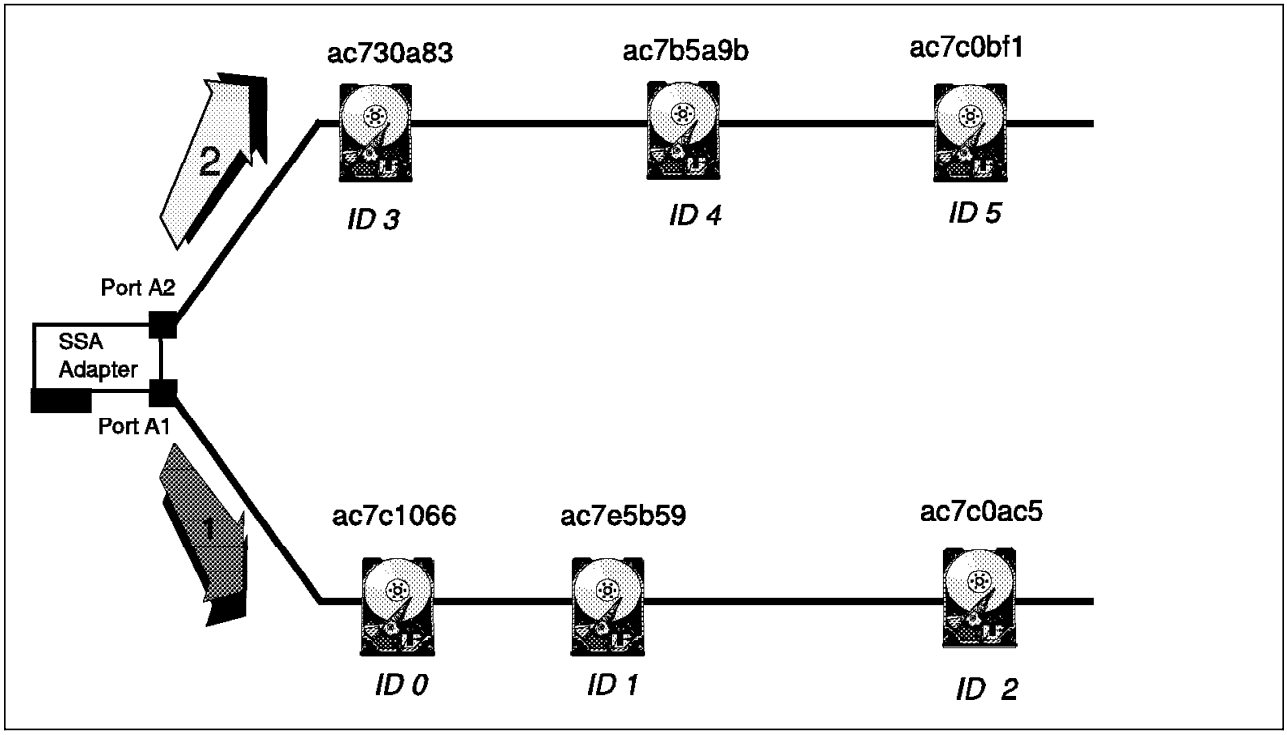

Figure 65. Logical Web Map - Two Strings

We recommend that you make the configuration persistent as soon as the system has detected and configured them in order to avoid SCSI IDs reassignment when the system is rebooted. The driver's configuration file,

/kernel/drv/sunligth.conf, maintains a persistent configuration database for each adapter.

Caution

When you install the driver for first time, the /kernel/drv/sunligth.conf file does not contain any databases. If you boot the system and the database for a given adapter does not exist, the default assignment rules are applied for that adapter.

4.1.2.5 Utilities Description

In this section we describe the utility programs and SSA services included with the IBM SSA Interface Controller for Sun SBus.

SSA Service Aid Utility: Use SSAU the SSA Service Aid Utility command **ssau**, to display the status of the SSA system and to diagnose problems with SSA disks, SSA adapters, and SSA system-level configurations.

The SSAU utility is automatically installed as part of the adapter software installation process. It uses a standard UNIX command line syntax so the programs are invoked from the shell prompt. To use SSAU, the operator must be logged on as root.

To see SSA Service Aid Utility command **ssau** help information, type **ssau -h**

The utility displays the text shown in Figure 66.

```
ssau help syntax

 ssau -h a/c/d/e

    a        adapter operation help

    c        health check operation help

    d       disk operation help

    e      error log analysis help
```

Figure 66.

In the sections that follow, we describe the functions of the SSAU utility.

Adapter Operations (option -a)

Use option -a when you want to display the adapter status or configure or test the adapter.*:*

To see a help file with a summary of the SSAU adapter operations syntax, type `ssau -h a`.

The utility displays the text shown in Figure 67 on page 80

```
       -a              specify adapter operation and adapter number

       -i              identify adapter(s)

                       optional specify duration of identify in seconds.
Default 20 sec.

 -v              display adapter vital product data

       -l "filename",(f)     download adapter f/w

                  ,f = force upgrade

       -t              Run adapter diagnostics (test)

       -w              Show web topology
```

Figure 67.

Adapter Specify **(option 0\1)***:* The SSAU utility assigns numbers 0 and 1 to the two adapters for the system.

For additional information refer to Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27.

Adapter Identify **(option -i)***:* This option is of great help if your system has more than one IBM SSA Interface Controller for Sun SBus adapter and you have to positively identify a particular adapter. The identify operation flashes the green LED of the specified adapter for a definable number of seconds; the default is 5 second.

For example,

ssau -a 0 -i identifies adapter 0 for 5 seconds.

ssau -a 1 -i 10 identifies adapter 1 for 10 seconds.

Adapter Vital Product Data **(option -v)***:* Use this option to obtain information about vital product data (VPD) such as a hardware or firmware revision or adapter memory size.

For example,

ssau -a 0 -v displays adapter 0 VPD

ssau -a 1 -v 10 displays adapter 1 VPD

Adapter Firmware Update **(option -l)***:* This option allows you to program the adapter flash memory with new firmware or with a version of firmware that is earlier than the version already on the adapter.

You must stop all disk accesses occurring through the adapter. Before running this option, one way to stop disk access would be to unmount the file systems on the affected disks.

The syntax of this command is:

- Update the firmware with the new version:

  ssau -a 0 -l (filename)

- Update the firmware with a version that is earlier than the version already on the adapter:

  ssau -a 0 -l (filename,f)

  The **,f** option forces the update.

Adapter Diagnostics Test **(option -t)***:* This option, which enables you to diagnose problems with the adapter, is of great usefulness because it executes the POST without power cycling the host system.

The command resets the adapter and reports the results of the POST.

You must stop all disk accesses occurring through the adapter.

The syntax of the command is:

ssau -a 0/1 -t

Show SSA Web Map **(option -w)***:* This option enables you to display the configuration of the SSA web attached to a given adapter.

The display shows the SSA UIDs of each node in the web as well the SCSI IDs for any nodes that are disk drives.

The syntax of the command is:

ssau -a 0/1 -w

For example, to display the SSA web configuration attached to adapter 0, enter:
`ssau -a 0 -w`

ssau Disk Operations **(option -d)***:* This option enables you to display disk status and configure and test the disk.

To see a summary of the ssau disk operations syntax, type:
`ssau -h d`

Figure 68 on page 82 shows the ssau disk operations.

```
          -d   (0/1) 0/1/2/.../47              specify disk
operation and drive number


optionally specify adapter number

            -i     (duration)                         identify
drive(s)


optional specify duration of identify in seconds. Default 20

                                                          sec.

            -v                                          display
drive vital product data (VPD)

            -l     "filename",(f)                   download drive

                                                        ,f =
force upgrade

             -t                                          run
drive diagnostics test

             -w                                         show web
topology

             -p

            (fd)                      restore factory settings

                 (sh)                      show current settings

                 (ac=on/off)                      set adaptive caching
on or off

                 (wc=on/off)              set write caching on or off

                 (rc=on/off)              set read caching on or off

                 (cs=1/2/4/8/16/32)       set the number of cache segments

              -s
set/reset service mode
```

Figure 68. SSAU Disk Operation

Disk Identify **(option -i)***:* When you execute this option, the yellow LED of the specified disk drive flashes for a definable number of seconds; the default duration is 20 s.

The syntax of the command is:

ssau -d (a:)n -i (duration) (identify disk **n** on adapter **a**)

For example,

ssau -d 6 -i identifies disk **6** on adapter **0** for 20 s.

ssau -d 1:6 -i 40 identifies disk **6** on adapter **1** for 40 s.

Disk Vital Product Data **(option -v)***:* Use this option to obtain information about the specified disk, such as hardware or firmware version.

The syntax of the command is:

ssau -d (a:)n -v display **vpd**of disk **n** on adapter **a**

For example,

ssau -d 6 -v displays VPD of disk **6** on adapter **0**

ssau -d 1:6 -v displays VPD of disk **6** on adapter **1**

Load Disk Firmware **(option -l)***:* This option enables you to update the firmware on the disk drives with new firmware or with a version of firmware that is earlier than the version already on the disk.

Before running this option, you must stop all data access to the disk whose firmware is being updated. For example, you could temporarily unmount the file systems on the disks affected.

The syntax of the command is:

ssau -d (a:)n -l *"*filename*"**(,**f)** (update firmware of disk **n** on adapter **a**) For example,

- Update the firmware with the new version:
  ssau -d 5 -l dchc85.bin (update firmware of disk **5** on adapter **0**)
- Update the firmware with a version that is earlier than the version already on the adapter,
  ssau -d 1:5 -l dchc81.bin,f (force update firmware of disk **5** on adapter **1**)
  The **,f** option forces the update.

Disk Diagnostics Test **(option -t)***:* This option enables you to diagnose disk drive problems. It causes the drive to execute its built-in test.

Before using this option, you must stop all data access to the disk.

The syntax of the command is:

ssau -d (a:)n -t (run diagnostic test on disk **n** on adapter **a**)

Set Disk Caching Parameters: This option enables you to display or set the caching parameters of a disk drive.

The syntax of the command is:

```
ssau -d (a:)n -p (set or display caching parameters
for disk n on adapter a)

        (fd)                restore factory settings

        (sh)                show current settings

        (ac=on/off)         set adaptive caching on or off

        (wc=on/off)         set write caching on or off

        (rc=on/off)         set read caching on or off

        (cs=1/2/4/8/16/32)  set the number of cache segments
```

Figure 69. Command Syntax

For example,

ssau -d 5 -p fs restores factory defaults to disk 5 on adapter 0.

ssau -d 1:5 -p ac=on sets adaptive caching on for disk 5 on adapter 1.

ssau -d 7 -p wc=on, cs=32 turns on write caching and sets number of cache segments to 32 for disk 7 on adapter 0.

ssau -d 1:0 -p sh shows current settings for disk 0 on adapter 1.

Set/Reset Drive Service Mode: This option isolates a disk drive, so that it can be replaced without interrupting the SSA loop.

The **set** option wraps the SSA ports of the nodes adjacent to the specified drive.

The **reset** option takes the adjacent nodes out of wrap mode after a replacement drive has been installed in the loop.

The syntax of the command is:

ssau -d (a:)n -s set/reset (set/reset service mode on disk n on adapter a)

For example:

ssau -d 9 -s set set service mode on disk 9 on adapter 0

ssau -d 1:20 -s reset reset service mode on disk 20 on adapter 1.

ssau Health Check: The utility checks the SBus adapter and the SSA disk drives attached to it. It performs the following functions:

1. Tests the firmware in the adapter to verify whether it is responsive to host driver commands
2. Obtains the current SSA web map
3. perform a SCSI inquiry command to each disk drive in the SSA web map to determine the redundant power and cooling status in the disk drives

The health check utility writes information in the Solaris system messages log (**/var/adm/messages**) if any of the above checks fails to complete or returns a negative status.

The utility can be run once, or it can be repeated at specified intervals.

The syntax of the command is: **ssau -c (interval) (repeat_count)**

where interval is the health check interval in minutes.

If you specify an interval, the system performs periodic health checks according to the interval. If you do not specify an interval, the system performs a single set of SSA health checks.

Repeat_count is the number of times to perform the health check.

If you specify an interval, the system performs the number of SSA health checks you specified in repeat_count.

Warning

If you specify an interval but not the repeat_count, the system performs the health check indefinitely in the interval specified.

For example,

SSAU -c 1 10 performs the health check function 10 times per minute

ssau Error Log Analysis: This utility enables you to analyze the Solaris system messages log (**/var/adm/messages**) for SSA-related messages that have been logged for the IBM SSA Interface Controller for Sun SBus driver or by the SSAU health check -d option.

The error log analysis function generates service request notification (SRN) messages based on its analysis of the logged SSA messages.The SRNs generated are logged in the SSA SRN log (/var/adm/ssasrn.log).

The syntax of the command is: **ssau -e (history) (interval)**

Where history is the amount of error log analysis to be processed in minutes.

The history parameter specifies how much history is to be considered by the error log analysis—in other words, how far back in the error log the analysis should go. If this parameter is not specified, the default time is 1440 min. (24 hr). The minimum allowed value is 15 min. The maximum allowed value is 9999 min.

If you specify the history parameter and omit the interval parameter, the error log analysis processes the past history of the error log and then monitors the error log for new entries. The new entries are processed immediately after they are appended to the error log. Error log analysis runs in this continuous processing mode for 1440 min. and then terminates.

If you specify both parameters or omit both parameters, the error log analysis processes the current contents of the error log and terminates.

The interval parameter is the analysis interval in minutes. It specifies how often a particular SRN will be reported for reoccurring conditions. The interval parameter value is also used in the processing of SRNs that are generated for multiple errors. If the interval parameter is not specified, the default time is 360 min. (6 hr). The minimum allowed value is 15 min. The maximum allowed value is 999 min. If the interval value is greater than the value of the specified history parameter, the interval is set to the value of the history parameter.

For example,

ssau -e performs the error analysis with default values and terminates.

SSA Configuration Utility: Use the SSA Configuration Utility (**SSCF**) to control the mapping of SSA to SCSI IDs.

The utility is automatically installed as part of the adapter software installation process. It uses a standard UNIX command line syntax, so the programs are invoked from the shell prompt. To use **sscf**, the operator must be logged on as root. The system administrator can use this utility to inspect, modify, and save the current mapping configuration as well as create or erase the persistent configuration database.

When you boot the system for the first time after the driver package is installed, the host adapter driver assigns a SCSI ID to each SSA device that it finds attached to the adapters. Thus the Solaris SCSI target driver can probe for all attached devices, which results in entries for those devices to be built in the /devices and /dev directories. These entries in turn allow applications to access the devices. The utility uses I/O control calls to communicate with the host adapter device driver. The I/O control calls enables the acquisition and modification of driver tables.

Command Line Syntax: Here is the command line syntax of the SSA Configuration Utility:

sscf (-a 0/1) (-c) (-d) (-e) (-h) (-m mapfile) (-p) (-r) (-u) (-y)

You can combine two or more options in any order.

Table 1 shows a description of the options.

Options descriptions

Table 1 (Page 1 of 2).

Option	Description
-a 0/1	Specifies the adapter that the utility will use. You must specify this option if the host has more than one adapter.
-c	Creates a new map for all SSA devices. This option forces assignments for all disk devices on the web. Refer to section 4.1.2.4, “Assigning Mapping Addresses” on page 75 for further information
-d	Displays the current SSA UID to SCSI ID configuration map

Table 1 (Page 2 of 2).	
Option	**Description**
-e	Erases the persistent configuration map. Useful only with the p option
Help	Displays expanded help syntax and utility version information
-m	Specifies the map file used to set the map
-p	Retains current configuration map for use when the system is rebooted.
-r	Notifies the driver that the SCSI ID assigned to the next drive removed will be reassigned to the next drive added. In other words, the SCSI ID is preserved
-u	Updates current map by assigning SCSI IDs to all known devices that do not yet have SCSI IDs assigned. The update option uses this procedure for assigning the SCSI IDs: The lowest unassigned SCSI ID is assigned to the first unassigned device; successively higher unassigned SCSI IDs are assigned to subsequent unassigned devices. All devices that have SCSI IDs assigned before the update retain those IDs
-y	Do not display confirming prompts for the -c, -m, -p, and -u options.

Using the Configuration Utility: Remember that root is the only user who can use the SSCF utility because it is restricted to the system administrator.

If you enter the command without options, a short-form usage message is displayed. Here is the command output:

usage sscf (-a 0/1) (-c) (-d) (-e) (-h) (-m mapfile) (-p) (-r) (-u) (-y)

If you enter the command with the -h option, a long-form usage message is displayed.

Figure 70 on page 88 shows the command output.

```
root>sscf -h

sscf, SSA Configuration Utility, version 1.7, create Dec 4 1997,
22:02:41


usage sscf (-a 0/1) (-c) (-d) (-e) (-h) (-m mapfile) (-p)
(-r) (-u) (-y)

         -a              specify adapter (0/1)

         -c              create new configuration map

         -d              displays adapter configuration map

   ;-e  erase, when is used
with -p

            -h                            display this help message

         -m                            set adapter's configuration
map from file

            -p                               set adapter's
configuration map across the boots

            -r                               replace device.
SCSI ID of the next device is reassigned to the

                                              subsequent device
added



            -u                               updateadapter's
configuration map with new devices

            -y                               do not display
prompts for -c, -m, -p and -u; force answer 'y'
```

Figure 70. Command Output from Long-Form

Displaying the Current Configuration Map: To display the current configuration map, you must use the -d option. Here is the command syntax:

usage sscf -a 0 -d

This command displays information about the configuration map in adapter 0.

Figure 71 on page 89 shows the output.

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID       SCSI

#High         Low       ID

0800          5aea40ca  00

0800          5aea44b4  01

0800          5aea442c  02

0800          5aea3dbe  03

0800          5aea4189  04

0800          5aea4145  05

0800          5aea448e  06

0800          5aea1610  07
```

Figure 71. Configuration Map Display

The list in Figure 71 shows the association of SCSI IDs 0 through 7 with SSA UIDs. The lines that contain a # character are comments. The SSA unique ID is identified by an 8-byte array of numbers that is unique to that device, the unique ID. The unique ID consist of three fields defined as shown in Table 2 Table 2 also shows the association of unique IDs with the first device shown in Figure 71

Unique ID Field Definitions

Table 2.

Bytes	Field Definition	Field in the list
0,1	Reserved = 0x0000	0x0000
2,3,4	IEEE organizationally unique identifier	00085a
5,6,7	Vendor-assigned product identifier	ea40ca

In Table 2 you can see that:

UID High take the positions 2 and 3 in the array. UID low take the positions 4,5,6,7 in the array. This is the value used as SSA UID.

Updating the Configuration Map: The -u option can be used to update the current map by assigning SCSI IDs to all known devices that do not yet have SCSI IDs assigned. The update option uses this procedure to assign the SCSI IDs (see Figure 72 on page 90): The lowest unassigned SCSI ID is assigned to the first unassigned device; successively higher unassigned SCSI IDs are assigned

to subsequent unassigned devices. All devices that have SCSI IDs assigned prior to the update retain those IDs. shows this procedure. Refer to 4.1.2.4, “Assigning Mapping Addresses” on page 75 for further information.

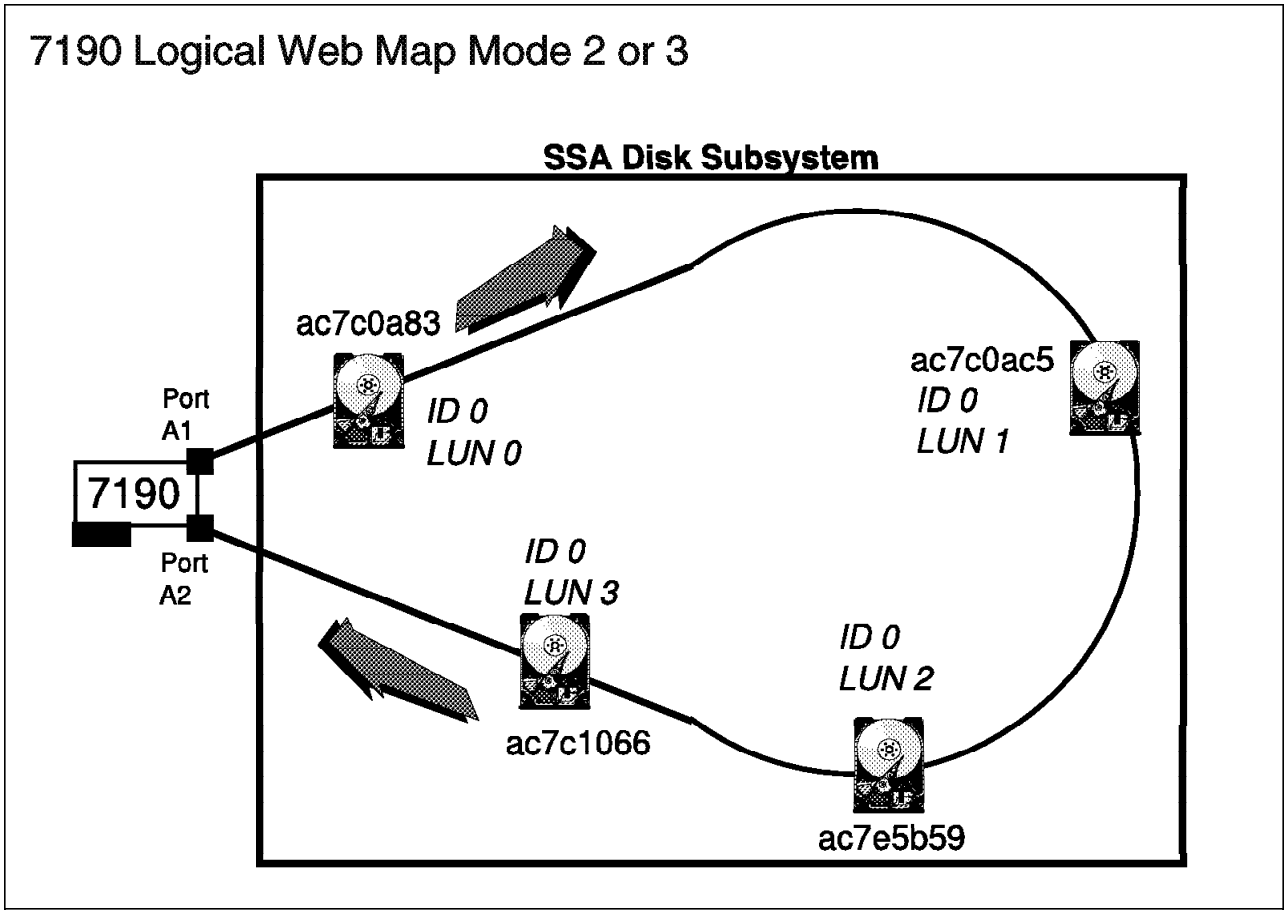

Figure 72. 7190 Logical Web Map Mode 2 or 3

After you create the new map, you must use the -p option to make the map persistent. For example, suppose the web on Adapter 0 consists of eight mapped devices. A group of eight new devices is then connected to the web. Because a map entry does not exist for these new SSA UIDs, the devices are not assigned SCSI IDs. To update the configuration map with assignments for these new devices, issue this command:

usage sscf -a 0 -u

The system displays the following information:

```
IDs of installed Adapters: 0 1

Continuing will create a default map for Adapter 0! Proceed?
(y/n) y

The configuration map has been modified.
```

To make the new configuration persistent across boots, use the **-p** option:

sscf -a 0 -p

The system displays the information shown in Figure 73.

```
IDs of installed Adapters: 0 1

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

Figure 73. Configuration file

Alternatively you can use both the **-u** and **-p** options:

sscf -a 0 -u -p

The system displays the information shown in Figure 74.

```
IDs of installed Adapters: 0 1

Continuing will update the map for Adapter 0! Proceed? (y/n) y

The configuration map has been modified.

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

Figure 74. Creating a New Configuration Map

The SSCF creates a new map for all SSA devices. The **-c** option makes the driver assign SCSI IDs to all disk devices known to be connected to the specified adapter. Be aware that the use of this option results in reassignment of SCSI IDs to SSA UIDs. Refer to 4.1.2.4, “Assigning Mapping Addresses” on page 75 for further information. This utility helps you define new devices that are attached while the system is live. For example, suppose the web on Adapter 0 consists of eight mapped devices, which are all then disconnected from the adapter. You connect a new group of eight devices to the adapter. Because a map entry does not exist for these new SSA UIDs, the devices are not assigned SCSI IDs.

You can list the last web map with the following command:

sscf -a 0 -d

figref refid=lawmap. shows the list of the last web map.

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID        SCSI

#High         Low        ID

0800          5aea40c1   00

0800          5aea44b2   01

0800          5aea4424   02

0800          5aea3dbf   03

0800          5aea4187   04

0800          5aea4146   05

0800          5aea448a   06

0800          5aea1622   07
```

Figure 75. Last Web Map Display

To create a new configuration map containing these new assignments, enter this command:

sscf -a 0 -c

The system displays the following information:

```
IDs of installed Adapters: 0 1

Continuing will create a default map for Adapter 0! Proceed?
(y/n) y

The configuration map has been modified.
```

You can list the new map with the following command:

sscf -a 0 -d

Figure 76 shows the list of the new map.

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID        SCSI

#High         Low        ID

0800          5aea40ca   00

0800          5aea44b4   01

0800          5aea442c   02

0800          5aea3dbe   03

0800          5aea4189   04

0800          5aea4145   05

0800          5aea448e   06

0800          5aea1610   07
```

Figure 76. New Web Map Display

To make the new configuration persistent across boots, use the option:

sscf -a 0 -p

The system displays the following information:

```
IDs of installed Adapters: 0 1

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

Alternatively you can combine the **-c** and **-p** options:

sscf -a 0 -c -p

The system displays the following information:

```
IDs of installed Adapters: 0 1

Continuing will create a default map for Adapter 0! Proceed?
(y/n) y

The configuration map has been modified.

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

Note

Upon first access to a newly powered up device, the Solaris target driver may display error messages because the device has not yet been spun up.

Using a Map File for Setting the Configuration Map

With option **-d** you can use a map file to specify the association of the SCSI IDs and SSA UIDs. Use of a map file provides more flexible mapping as well as the ability to recall and apply previously saved mappings.

Warning.

The use of the -d option results in reassignment of SCSI IDs to SSA UIDs. In other words, the map file associations override the current configuration map. Refer to 4.1.2.4, “Assigning Mapping Addresses” on page 75 for further information.

A map file is a plain ASCII file, so you can use a standard text editor to create or modify it. You can generate a file containing the state of the current mapping, redirecting the output of the utility′s -d option.

For example, to capture the mapping for Adapter 0 in file mapfile, enter:

sscf -a 0 -d > mapfile

Figure 77 shows the contents of the map file.

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID        SCSI

#High         Low        ID

0800          5aea40ca   00

0800          5aea44b4   01

0800          5aea442c   02

0800          5aea3dbe   03

0800          5aea4189   04

0800          5aea4145   05

0800          5aea448e   06

0800          5aea1610   07
```

Figure 77. Map Contents

The list in Figure 77 shows the association of SCSI IDs 0 through 7 with SSA UIDs. The lines that contain a # character are comments. Suppose that you want to reverse the association of all SCSI IDs and UIDs. Modify the map file, using a text editor, so that its contents are as shown in Figure 78 on page 96

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID       SCSI

#High         Low       ID

0800          5aea40ca  07

0800          5aea44b4  06

0800          5aea442c  05

0800          5aea3dbe  04

0800          5aea4189  03

0800          5aea4145  02

0800          5aea448e  01

0800          5aea1610  00
```

Figure 78. Modified Map

After you ensure have that I/O to all disk devices has ceased, use the **-m** option to reconfigure the mapping:

sscf -a 0 -m mapfile

The system displays the following information:

```
IDs of installed Adapters: 0 1

Continuing will apply the map for Adapter 0! Proceed? (y/n) y

The configuration map has been modified.
```

To make the new configuration persistent across restarts, use the **-p** option:

sscf -a 0 -p

The system displays the following information:

```
IDs of installed Adapters: 0 1

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

Alternatively, you can combine the **-c** and **-p** options:

sscf -a 0 -p -m mapfile

The system displays the following information:

```
IDs of installed Adapters: 0 1

Continuing will apply the map for Adapter 0! Proceed? (y/n) y

The configuration map has been modified.

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

4.1.2.6 Drive Configuration

SSA drives are automatically configured in the SSA loop, and the SSA adapter card in your system recognizes the configured drives. One additional step is required for each adapter card to save the drive configuration data. This section describes the process of configuring the drives in the server for the first time.

Enter the following command at the shell prompt, to make the new configuration persistent across boots:

sscf -a 0 -p

The system shows the following information:

```
IDs of installed Adapters: 0 1

About to modify file '/kernel/drv/sunligth.conf' Proceed? (y/n) y

Backup file '/kernel/drv/sunligth.conf.bak' has been written.

The driver's configuration file '/kernel/drv/sunligth.conf' has
been modified.
```

To configure the drives on a second Adapter card, enter the following command:

sscf -a 1 -p

Your new SSA drives are now assigned permanent SCSI IDs, which are saved during power-off cycles. The system always recognizes and assigns the same SCSI IDs even if you disconnect and reconnect these drives. To display the current configuration map that shows the SSA UIDs of your disk drives and their association with the SCSI IDs, use the following command:

usage sscf -a 0 -d for adapter 0 or

usage sscf -a 1 -d for adapter 1

Figure 79 on page 99 shows the output.

```
#IDs of installed Adapters: 0 1

#Configuration map for Adapter 0

#UID          UID       SCSI

#High         Low       ID

0800          5aea40ca  00

0800          5aea44b4  01

0800          5aea442c  02

0800          5aea3dbe  03

0800          5aea4189  04

0800          5aea4145  05

0800          5aea448e  06

0800          5aea1610  07
```

Figure 79. Output from New Command

We recommend that you print and save for future reference the information in the current configuration map. Use the following command syntax:

usage sscf -a 0 -d > myfile0 for adapter 0 or

usage sscf -a 1 -d > myfile1 for adapter 1.

where myfile is the result file to which the command output is redirected.

4.1.2.7 Configuring Multi-initiator Clusters

When you configure multi-initiator clusters, multiple host systems and storage devices are interconnected through SSA. Each host has direct access to all of the storage devices because the SSA loop configuration provides the connectivity. This feature allows storage devices to be shared between servers. If you want shared device access, you must provide volume-sharing software to coordinate data access. The volume-sharing software is required to prevent the potential loss or corruption of data.

4.1.2.8 Cluster Configurations

A cluster can have storage failover capability or concurrent shared storage if the correct disk storage software is installed. See Figure 80 on page 100

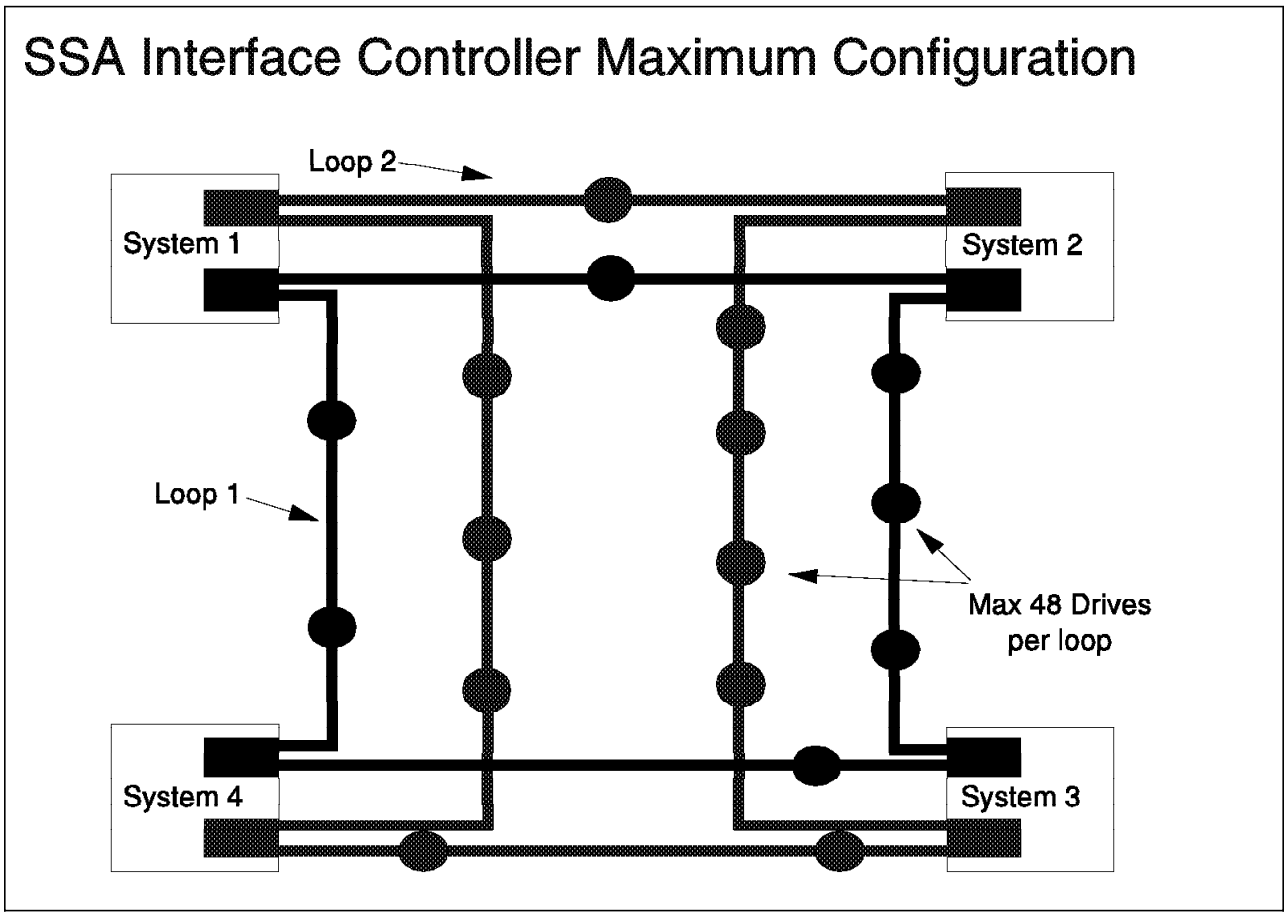

Figure 80. Cluster Configuration

If you use cluster with storage failover capability, each server in the cluster is assigned a set of storage devices that it does not share. If the server fails, its storage is reassigned by the other cluster management software to another working server. If you use clusters with concurrent shared storage, each server in the cluster has access to shared storage. The cluster management software provides the concurrent access controls.

SSA Configuration: The SSA topology requirements for multi-initiator clusters depend on the cluster management software. The storage device configuration requirements also depend on this software. The instructions for configuring a multi-initiator cluster are provided with the management software.

IBM requires the use of 9.1 GB disk drives to ensure maximum flexibility for setting up multi-initiator cluster configurations.

In cluster configurations, it is important that all systems in a cluster see a consistent view of the connected drives. This can be accomplished by using the **-m** option of the SSA Configuration Utility (see “Using a Map File for Setting the Configuration Map” on page 94.). The utility allows the generation of the map file based on the SCSI ID to SSA UID associations as viewed by one system. This map file must be propagated and applied to all other systems in the cluster. The result is a consistent view of the drives by all systems in the cluster.

To create the map file, use this command:

sscf -d > mapfile.cluster

Transfer the map file to other machines in the cluster and apply the map file, use this command:

sscf -p -m mapfile.cluster

4.2 IBM 7190 model 100

The IBM 7190 Model 100 provides the benefits of IBM's SSA Disk Subsystem on host systems that do not have a native SSA interface. See Figure 81 The SSA connection to the host system is through the host's SCSI bus. Because SSA devices appear SCSI like devices to the host system, installation and operation is simplified. The IBM 7190 Model 100 is a SCSI target device that creates and controls the connection to an SSA serial loop. We recommend that the 7190 Model 100 be the only SCSI device on the SCSI bus.

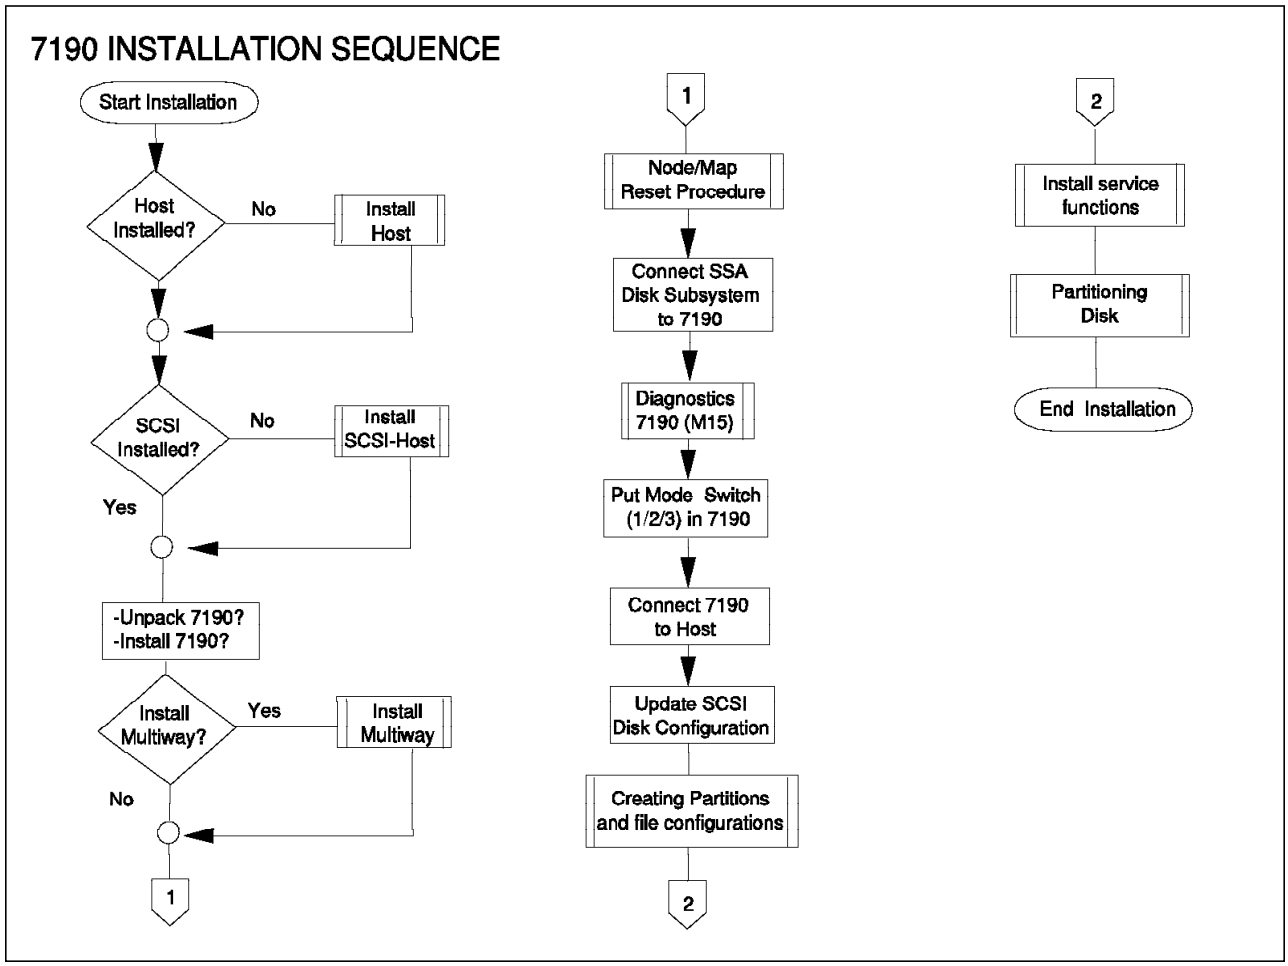

Figure 81. 7190 Installation Sequence

For further information, refer to Chapter 4, “Configuration and Installation” on page 65 and Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27. In this section we discuss how to prepare for installation and configuration of the IBM 7190 Model 100 in your system.

4.2.1 Preparing for Installation

The procedures in this section help will you prepare to install and verify the operating conditions for the IBM 7190. The following topics are covered:

1. Tools and information
2. Checking the inventory
3. Mode selection switch settings

4.2.1.1 Tools and Information

Before you start the installation, ensure that you have the following tools, information, documentation, and Web site.

Tools

You may need a small flat-head screwdriver for securing the SSA connectors and the SCSI to the 7190. It is possible to use the special screwdriver supplied with the 713x to tighten the SSA connector screws.

Information

Obtain the following information from your system administrator:

- Host system and SCSI available to connect the 7190 Model 100
- Host system shutdown and startup procedure

Publications and web site

You may also want to look at these reference manuals:

- *7133 SSA Disk Subsystem for Open Attachment: Installation and User's Guide, SA33-3273*
- *7133 SSA Disk Subsystem for Open Attachment: Service Guide, SY33-0191.*
- *Translated Safety Notices for Open Attachment: GC26-7246.*
- *7131 Model 405, SSA MultiStorage Tower for Open Attachment: Service Guide, SY32-0405.*
- *7131 Model 405, SSA MultiStorage Tower for Open Attachment: Installation and User's Guide, SA26-2001.*
- Addition information is available on the SUPPORT pages at the IBM storage home page:

  http://www.storage.ibm.com

4.2.1.2 Checking the Inventory

Verify that all items in your order were received, If you detect shipping damage or missing items, do not install the 7190 Model 100 without IBM approval. Report all observed damage or missing items to your IBM service representative.

4.2.1.3 Mode Selection Switch Settings

The 7190 Model 100 has a rotary push-button switch for selecting addressing or diagnostic modes. The mode selection switch is located on the right-hand side near the rear of the 7190 Model 100 (Figure 82).

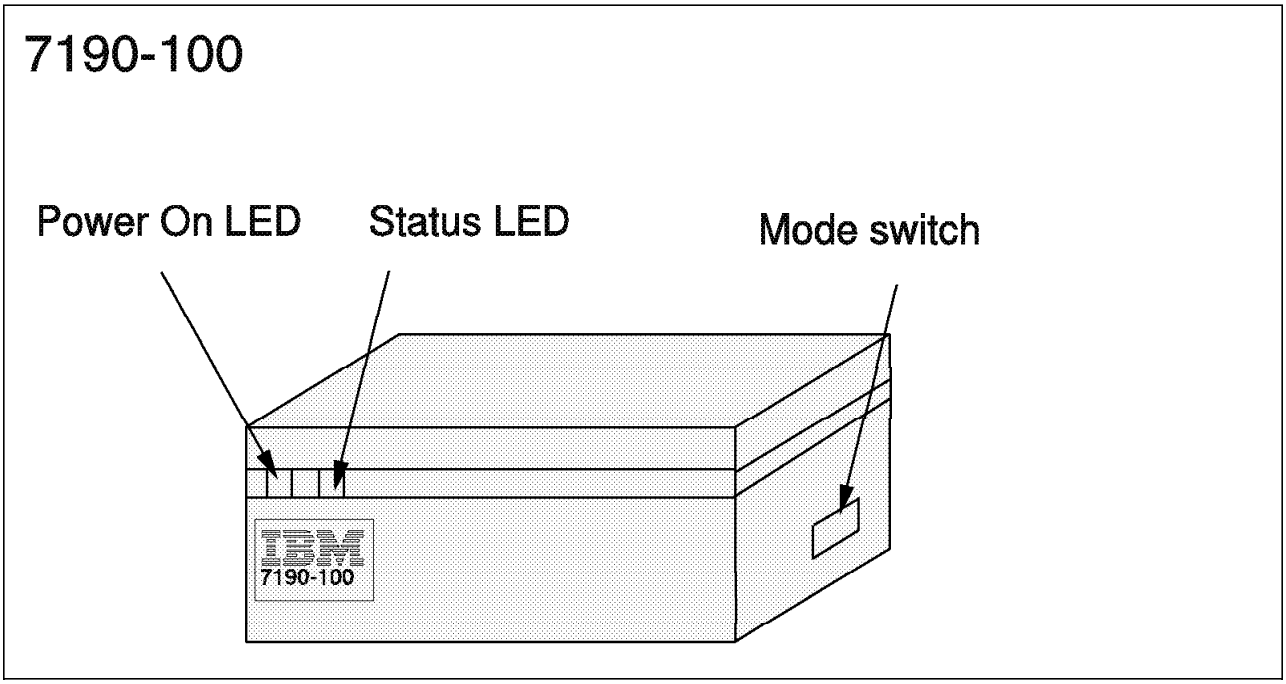

Figure 82. The 7190 Model 100 LEDs and Mode Switch

You can change the switch mode by pressing either black button, to increase or decrease the mode selection number. Table 3 summarizes SCSI addressing and diagnostic modes.

Addressing and diagnostic modes

Table 3 (Page 1 of 2).

Mode	Description
0	ID priority mapping, Fast/Wide SCSI, with device spin up. Uses SCSI IDs 0 trough 15 except ID7. This mode attempts to power up the drive motors automatically. Allow about two minutes for the drives come up to full speed before rebooting the host system.
1	Reserved
2	LUN priority mapping, Fast/Wide SCSI, Uses SCSI IDs o trough 15, except ID6 and ID7

Table 3 (Page 2 of 2).

Mode	Description
3	LUN priority mapping, Fast/Wide SCSI, with Device spin up.This mode attempts to power up the drive motors automatically. Allow about two minutes for the drives come up to full speed before you rebooting the host system.
4-12	Reserved
13	7190 microcode version level is presented on the status LED
14	Reserved
15	Subsystem Diagnostic Test Mode

Remember

When you change the switch mode, for the switch settings to take effect, you must power off the 7190 Model 100 and then power it back on.

4.2.2 Installation Procedure

In this section we explain the procedure to install the IBM 7190. We cover the following topics:

1. Safety information
2. 7190 Model 100 installation
3. Check installation errors
4. Assigning mapping address
5. Software installation
6. Utilities description

4.2.2.1 Safety Information

Lethal voltages may be present in your system even with the power switch off. The following safety notices are used in this installation procedure:

- Please follow all electrical safety instructions described in your system documentation.
- An electrical outlet that is not correctly wired could place hazardous voltage on metal parts of the system or the devices that attach to the system. It is the responsibility of the customer to ensure that the outlet is correctly wired and grounded to prevent an electrical shock.
- Before installing or removing parts, ensure that the power cables for the system unit and all attached devices are unplugged. Electrostatic discharge (ESD) protection is required for your system at all times.

Translations of these notices are in *Translated Safety Notices for Open Attachment, GC26-7246*

4.2.2.2 7190 Model 100 Installation

In this section, we describe the procedure for installing the 7190. This procedure assumes that your host system is properly installed and operating according to the host system manuals. You must ensure that the system hardware and system software meet the requirements defined in Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27. Figure 83 is an overview of the steps to install the 7190.

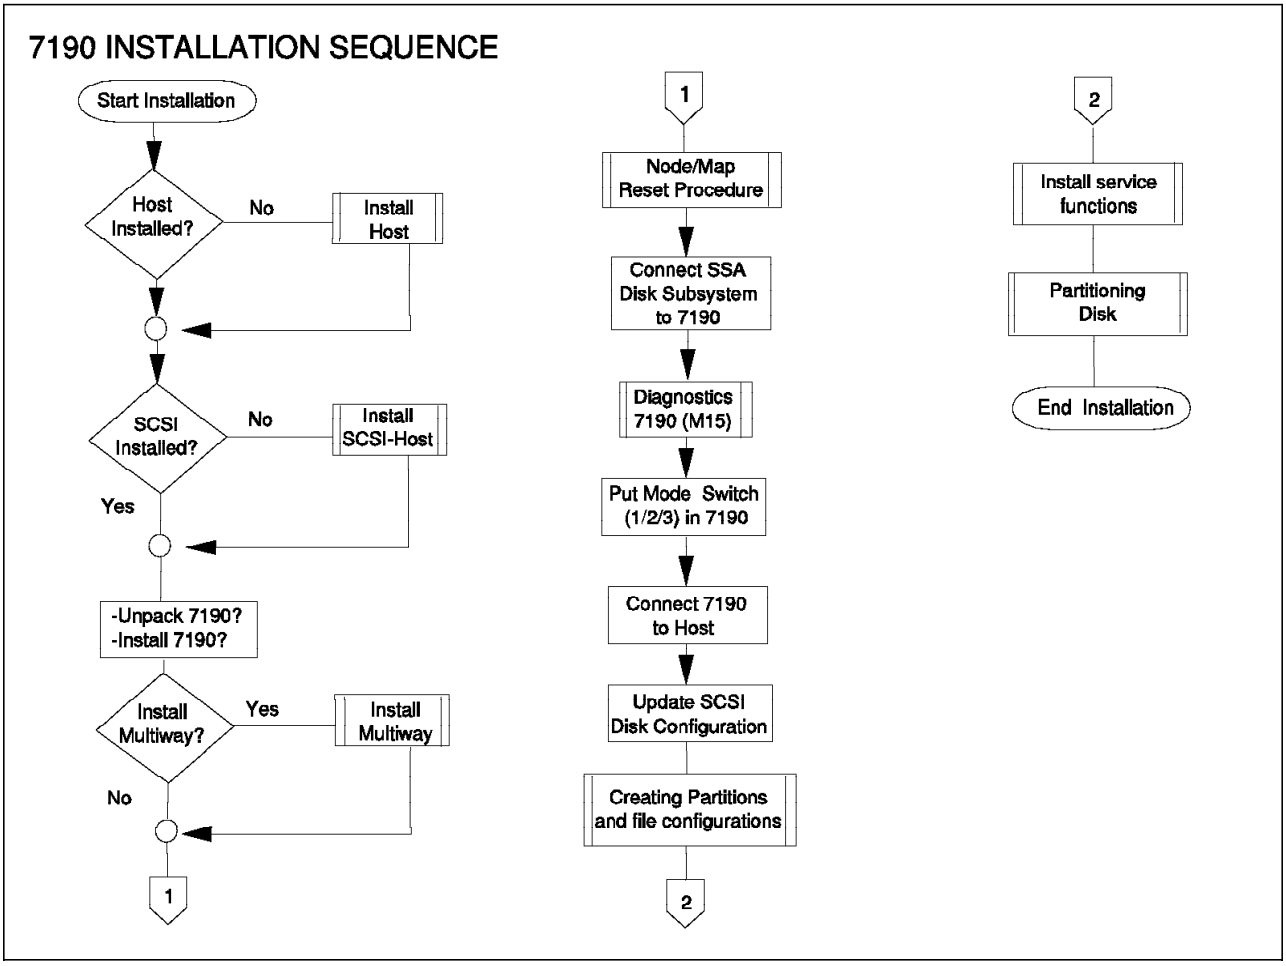

Figure 83. IBM 7190 Model 100 Installation Sequence

We recommend that you back up or copy all files before making major changes to the system.

You are ready to install the 7190. Follow these steps:

1. Place the 7190 model 100 in a location that is convenient for your system and remember that a minimum of 25.4 mm or 1 in. of open space adjacent to the unit is required for ventilation.
2. If you are installing a multiway configuration, verify that the 7190 model 100 is at level 6.3 or above. The following is the procedure for verifying the microcode level in the 7190:

Microcode Revision Level Display: Press the push-button switch (Figure 84 on page 106) until the number 13 appears. This mode allows you to display the microcode revision level. The 7190 Model 100 displays the level using a flashing

code in the status LED. See the following status code pattern example to read the revision level.

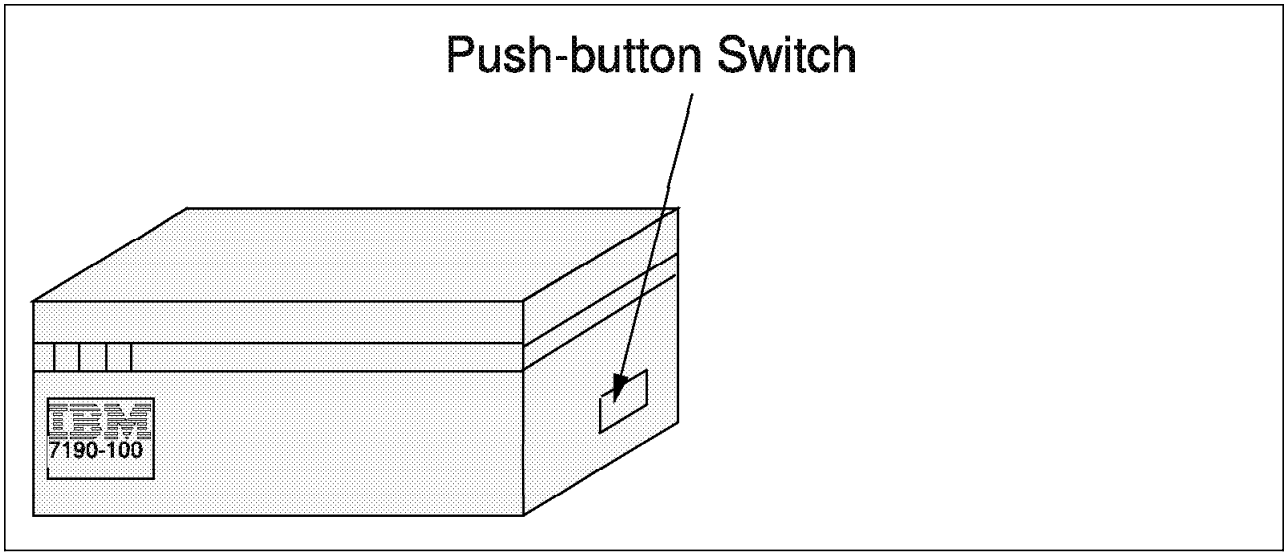

Figure 84.

You can identify the flashing status code, observing the pattern and duration of pauses and the duration of LED flashing. The decimal digits are presented in the following manner:

0=short fast blink

1=LED on once

2=LED on twice

3=LED on three times

and so on.

For example, the status code 12, 3, 0 has the following pattern:

Long pause (LED off)

Flashing decimal number, LED on once=decimal 1

Short pause (LED off)

Flashing decimal number, LED on twice=decimal 2

Medium pause (LED off)

Flashing decimal number, LED on three times=decimal 3

Medium pause (LED off)

Flashing decimal number, LED short fast blink=decimal 0

Long pause (LED off)

The pattern repeats itself.

A long pause separates the repeated presentation of the status, and a medium-length pause separates the digits of the code.

Node Map Reset Procedure: The 7190 Model 100 keep informations about the web map into nonvolatile RAM. This procedure must be carried out to be sure that the 7190′s nonvolatile RAM does not contain information when the 7190 Model 100 begins the web talk and assign SCSI IDs to devices in the SSA disk storage subsystem.

You are ready to perform the node map reset for the 7190. Follow these steps:

1. Set the power switch to off.
2. Install an SSA cable directly between the A1 and A2 connectors on the 7190 model 100 enclosure (see Figure 85). Ensure that there are no devices on the SSA loop.

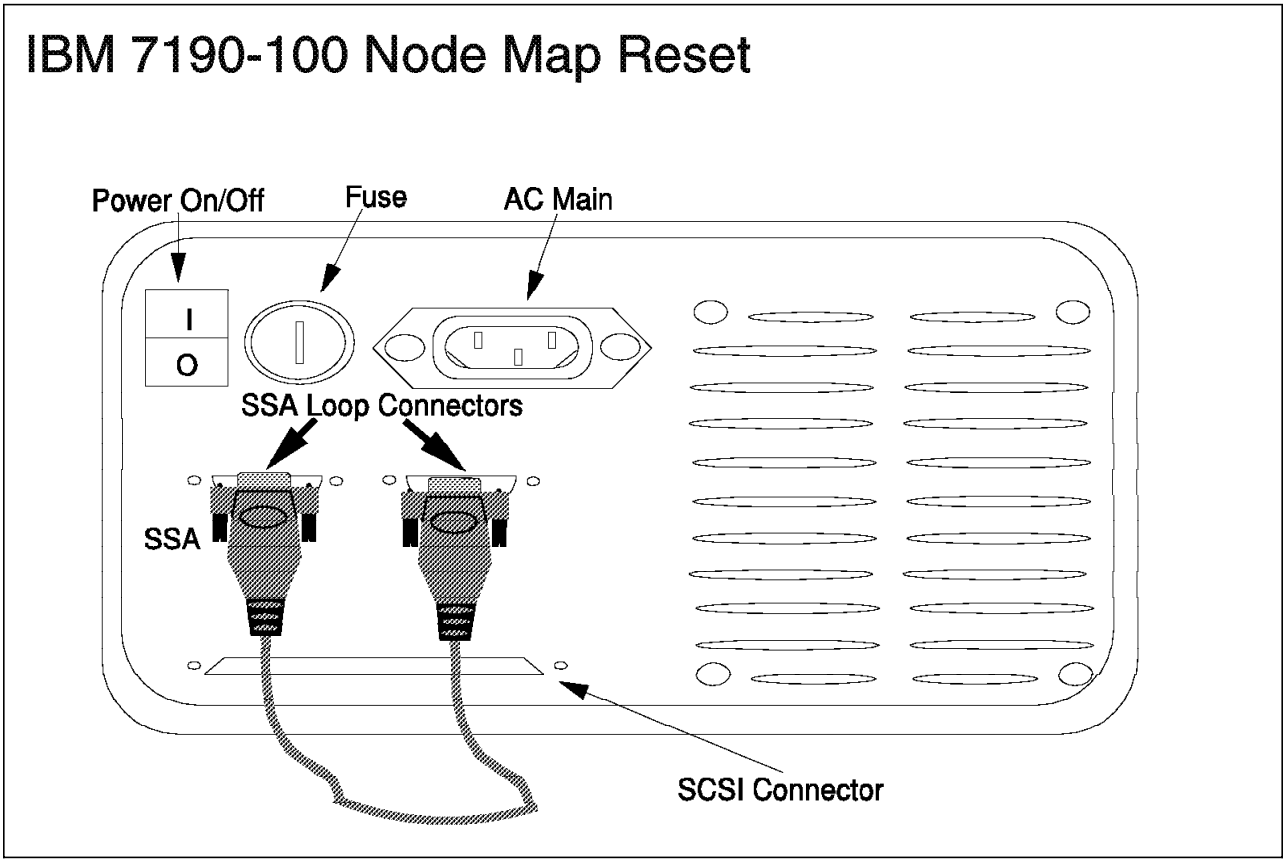

Figure 85. The 7190 Model 100 Cable Connection for Node Map Reset

3. Set the power switch to on.

   Wait for the activity LED to blink status code 060. The node mapping table into the 7190 nonvolatile RAM is now cleared of the all previous entries.

4. Set the power switch to off.
5. Remove the SSA cable between the A1 and A2 connectors.

Now you are ready to install the 713x SSA Disk Storage Subsystem.

Installing the 713x SSA Disk Storage Subsystem: An IBM service representative may be required to install the 713x SSA disk storage subsystem. Follow these steps:

1. Turn the 713x SSA disk storage subsystem unit off.
2. Check the power cables and make sure that they are removed from the power outlet.
3. Install the SSA cables between the 7190 model 100 and the SSA disk storage subsystem.

Figure 86 shows this procedure.

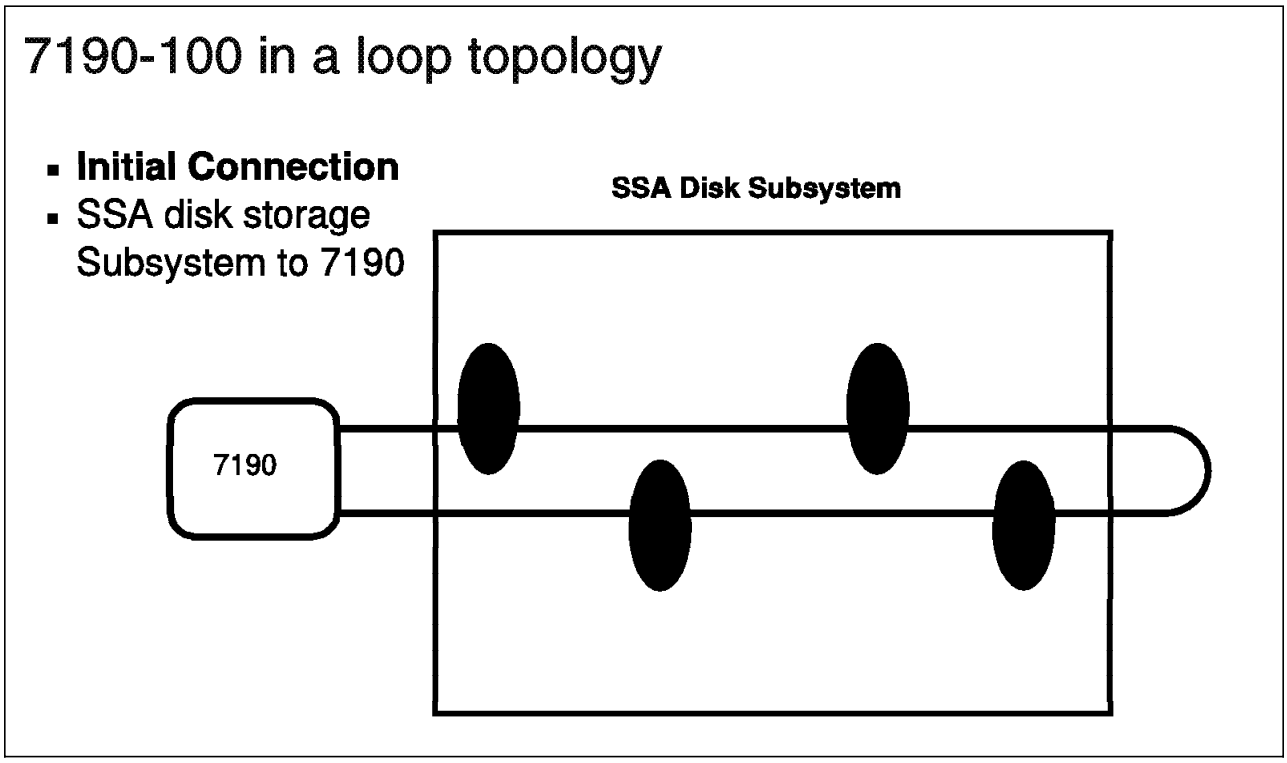

Figure 86. IBM 7190 Model 100 and SSA Disk Storage Subsystem Connection

Remember to make an SSA loop. You need at least two external SSA cables and at least one SSA device. Connect one end of one cable to either port on the 7190 Model 100 and the other end to one of the ports on the storage device. Then connect one end of the other cable to the other port on the 7190 and the other cable end to the other port on the storage device. Multiple devices can be added to the loop. For each device added, you need additional external SSA cables. After you have set up the devices and cables in the configuration you want and according to Chapter 3, “IBM 7190 and SBus Adapter Planning and Design” on page 27, make sure the cables are tightly screwed to the device and the 7190 Model 100 ports. Select mode 15 on the 7190. (See Figure 87 on page 109.) This mode enables you allows to perform the subsystem diagnostics test and test the 7190 Model 100. Proceed to verify the disk drives over the SSA loop. See Figure 87 on page 109.

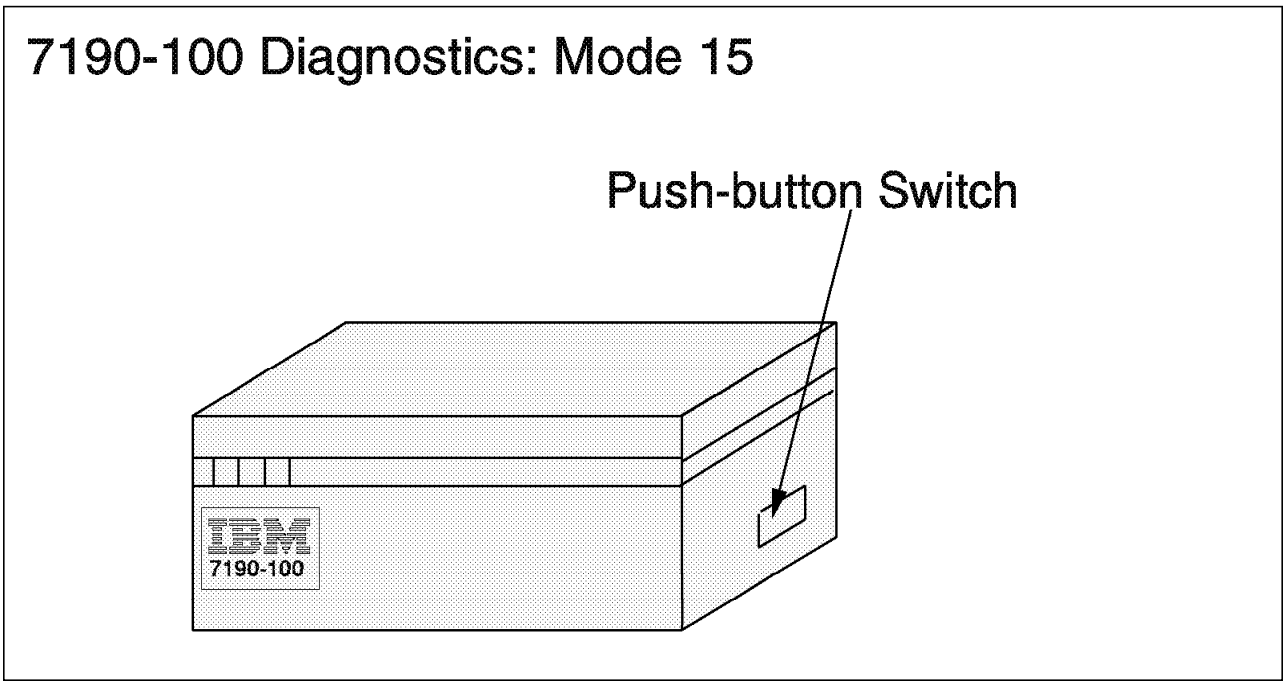

Figure 87. IBM 7190 Model 100 Diagnostics Mode 15.

Now are you ready to power up the SSA devices and your 7190.

Here is the power-up sequence that we recommend:

1. Power up the 713x SSA storage disk subsystem Set the 713x storage disk subsystem power switch to on.
2. Power up the 7190 Model 100

Use the following procedure to verify the status of the devices:

1. Check the LEDs on all drives in the 713x SSA storage disk subsystem.

   The green LEDs in each drive should be on and blinking. When all the SSA devices pass the test, the blinking drive LEDs are turned off because the 7190 model 100 resets all drives.

2. Check the status LED on the 7190 Model 100.

   If the diagnostic test detects an error, the status LED blinks an error code. You must proceed to identify the error code. The code identified must have three digits.

   - The first number specifies the device that failed in the loop.
   - The second number specifies the SSA LUN of the device that failed in the loop. For SSA devices, this number should always be 0.
   - The Third number shows the specific error.

   Table 4 on page 110 lists the diagnostic error codes and describes them.

Diagnostic Error Description

Table 4.

Status Code	Error Description
001	ID No inquiry data received.
003,003	Wrong inquiry data
004	Inquire command time out
005	Test unit ready time out
006	Unexpected sense data/poor SCSI connection. Check the SCSI cable
007	Stop unit time out
008,009,010	Drive spin up error
011,012,013,014,020	Read capacity command error
015,021	Read command error
016,017,019,023	Not recovered
018,024,025	Write command error
022	Recovered error
026	Mode selection command error
090	No response from drive
091,092,093	Sense command error
other codes	Possible attachment error

3. Once you are sure all of the drives have been tested and there are no errors, set the 7190 Model 100 switch to off. Set the mode switch to the SCSI setting appropriate to your host system requirements.

Now are you ready to connect the 7190 Model 100 to host.

Connecting the 7190 Model 100 to the host: Set the host system power switch to off. Connect the shielded SCSI cable to the host system according to the procedures in the system documentation. Connect the shielded SCSI cable to the 7190 model 100 SCSI connector.

The Figure 88 on page 111 shows the procedure.

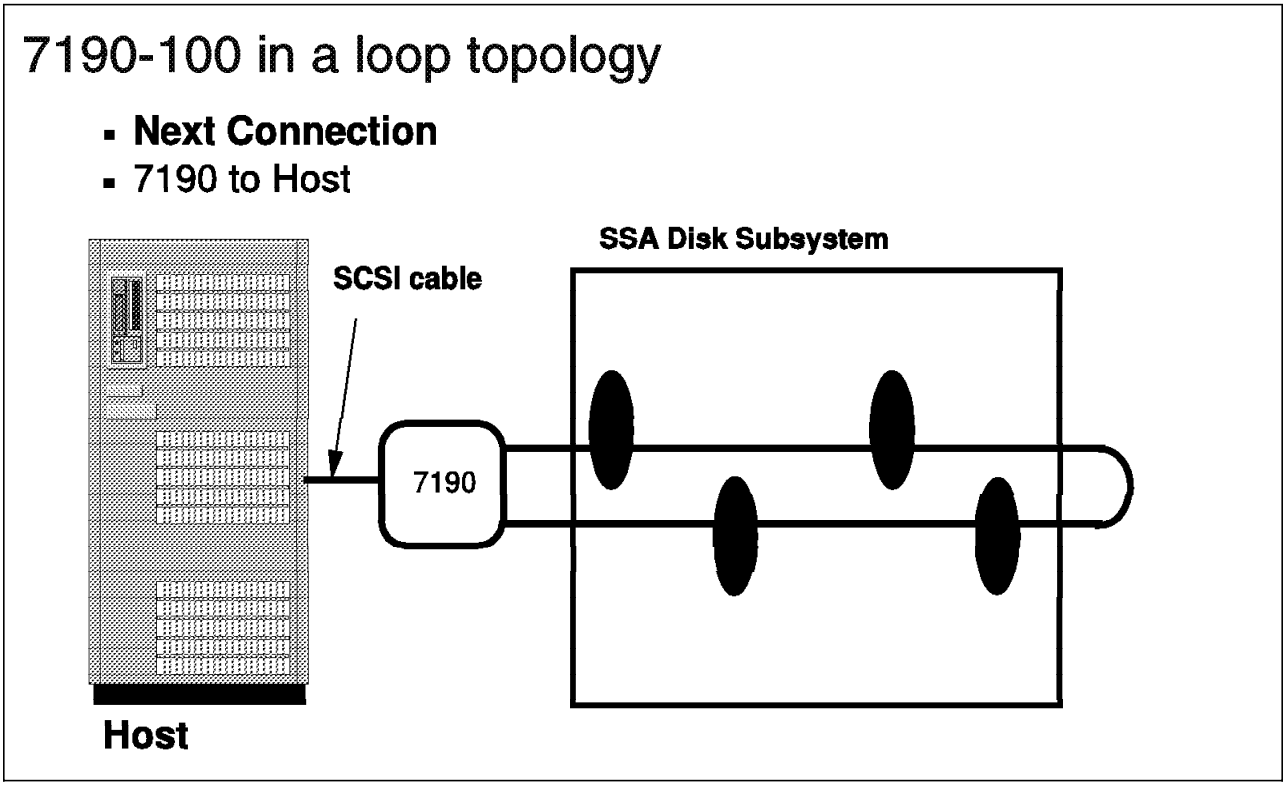

Figure 88. IBM 7190 Model 100 and host connection .

Now are you ready to power up the host and the 7190.

Here is the power-up sequence that we recommend:

1. Power up the 7190 model 100. The status LED is off when the power switch is set to on. The LED turns on and remains on when at least one SSA port becomes active.
2. Power up the host.

If error conditions occur while you are installing the 7190 Model 100, then refers to " Check Installation Errors."

Check Installation Errors: In this section we present the step to follow if problems occur during the installation of the 7190:

1. If the status LED is off, This indicates that the 7190 Model 100 cannot find any SSA devices. We recommends that you:
   - Check whether the SSA cables are properly connected to the 7190 model 100 and the SSA devices.
   - Check whether there is power to your SSA devices.
   - Check the SCSI cable and replace if necessary.
   - Check the SSA devices by attaching them to a 7190 Model 100 on another host system with an SSA interface to ensure that the devices are working correctly.
2. If the status LED is in a fast blinking pattern, The 7190 Model 100 is remapping the devices on the web or a download microcode is in progress.

3. If the status LED is flashing in a 90% on and 10% off pattern, the SSA loop is open and the devices are configured as two strings.
4. If the status LED is on but not flashing, the SSA ports are connected in loop.
5. If the status LED is flashing in an intermittent pattern, the 7190 model 100 is in a fault condition and is unable to continue the operation. Verify the status code in the 7190.

   The most frequent cause of this condition is related to SCSI or SSA cabling.

If you still cannot resolve the error situation after you have checked the installation errors, refer to Chapter 5, “Maintenance” on page 125 for further information about problem determination.

4.2.2.3 Assigning Mapping Address

A UID identifies an SSA device, and a SCSI ID identifies a SCSI device. The operating system recognizes SSA disks as SCSI disks because a mechanism that maps SCSI IDs to SSA UIDs. The 7190 Model 100 maps attached SSA devices to SCSI/LUN pairs. This section describes how SCSI ID/LUN pairs are mapped to SSA devices on the SCSI bus.

The 7190 Model 100 usually responds to SCSI IDs 0 through 5 and 8 through 15. If you use mode 0, the 7190 Model 100 also responds to SCSI ID 6. The 7190 Model 100 contains a SCSI ID/LUN-to-node mapping table in nonvolatile RAM. When the 7190 Model 100 is powered on, it scans the SSA bus. If detects a drive that is not in the node mapping table, it makes a new entry containing the new drive ID and the next SCSI ID/LUN pair. Here is the order in which the 7190 Model 100 assigns SCSI ID/LUN pairs:

The 7190 Model 100 scans disk drives on the SSA bus. The first drive is the drive closest to the A1 SSA port, and the second drive is the SSA drive next closest to the A1 port. This process continues until the farthest drive from the A1 port. If you connected the drives in a loop, the last drive of the scan is the drive closest to the A2 port. (Figure 89 on page 114.) If you connected drives in string (the ends of the loop are open), the ordering continues with the drive closest to the A2 port. The scan ends with the drive farthest from the A2 port. (Figure 90 on page 115.) The process for assigning SCSI ID/LUN pairs uses mode 3 in the 7190 Model 100, which means LUN priority.

SCSI Address Mode Selection in the 7190 Model 100: This subsection describes which SCSI IDs and LUNs the 7190 Model 100 uses when mapping SSA devices to SCSI bus. A rotary switch on the side of the 7190 allows you to select either of the two addressing mode defined in the 7190. The 7190 Model 100 is shipped with the mode selection switch set to mode 3.

Mode Selection: The 7190 has a rotary push-button switch for selecting addressing or diagnostic modes. The mode selection switch is located on the right-hand side near the rear of the 7190.

You can change the switch mode by pressing either black push-button to increase or decrease the mode selection number. See Table 5 on page 113 for the SCSI addressing and diagnostic modes.

Addressing and Diagnostic Modes

Table 5.	
Mode	**Description**
0	ID priority mapping, Fast/Wide SCSI, with device spin up. Uses SCSI IDs 0 through 15 except ID7. This mode attempts to power up the drive motors automatically. Allow about 2 minutes for the drives to come up to full speed before rebooting the host system.
1	Reserved
2	LUN priority mapping, Fast/Wide SCSI, Uses SCSI IDs 0 through 15, except ID6 and ID7
3	LUN priority mapping, Fast/Wide SCSI, with device spin up.This mode attempts to power up the drive motors automatically. Allow about 2 minutes for the drives to come up to full speed before you reboot the host system.
4-12	Reserved
13	7190 microcode version level is presented on the status LED
14	Reserved
15	Subsystem diagnostic test mode

Mode 0 is defined as ID priority, Fast/Wide, with device spin. When mode 0 is selected, the 7190 Model 100 uses ID priority to map SSA drives to SCSI ID/LUN pairs. When you use ID priority mapping, the first SCSI ID/LUN pair is ID0 LUN0, followed by ID1 LUN0, ID2 LUN0, continuing to ID6 LUN0. The numbering then skips to ID8 LUN0 continuing until ID15 LUN0. The sequence is repeated, increasing the LUN by one. This method provides 120 SCSI ID/LUN pairs.Table 6 shows mode 0 ID priority. Figure 89 on page 114 shows the web map assigned sequence in a loop connection, and the Figure 90 on page 115 shows the web map assigned sequence in a string connection.

Mode 0 - ID Priority

Table 6 (Page 1 of 2).	
ID	**LUNs**
ID0	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID1	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID2	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID3	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID4	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID5	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID6	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID8	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID9	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID10	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7

Table 6 (Page 2 of 2).	
ID	**LUNs**
ID11	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID12	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID13	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID14	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID15	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7

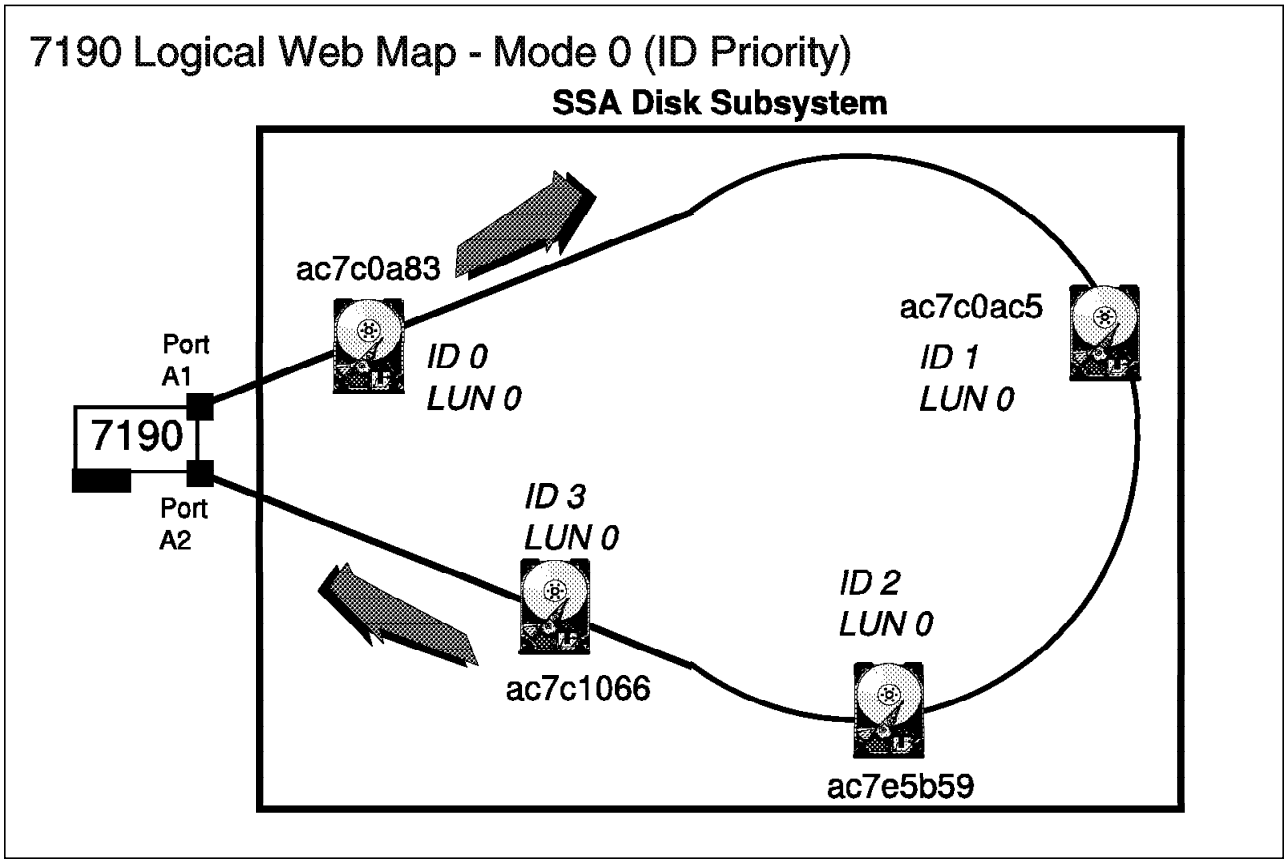

Figure 89. IBM 7190 Model 100 Mode 0 Loop Connection

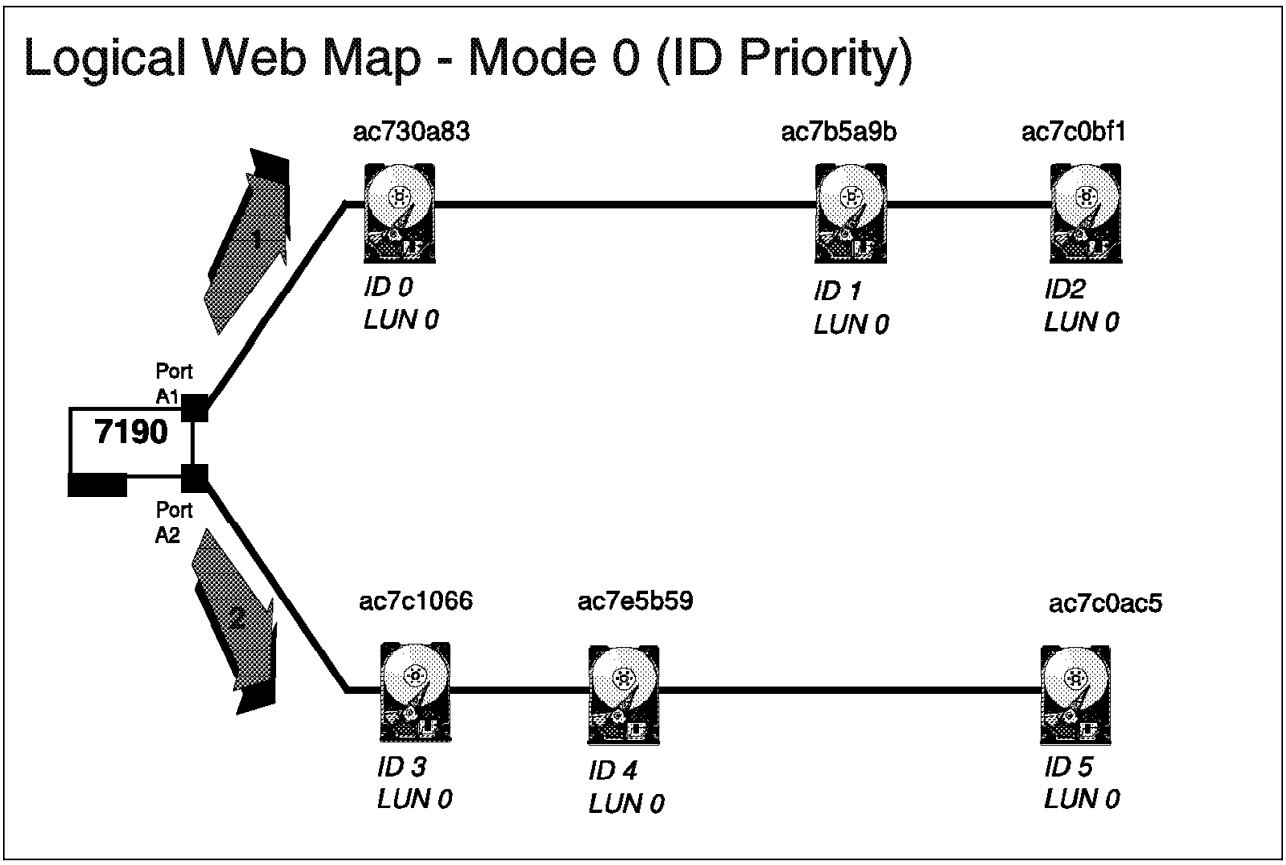

Figure 90. IBM 7190 Model 100 Mode 0 String Connection

Mode 2 is defined as LUN priority mapping, Fast/Wide SCSI. When mode 2 is selected, the 7190 Model 100 uses LUN priority to map SSA drives to SCSI ID/LUN pairs. When you use LUN priority mapping, the first SCSI ID/LUN pair is ID0 LUN0, followed by ID0 LUN1, ID0 LUN2, continuing to ID0 LUN7. The sequence is repeated, increasing the ID by one. This method does not assign ID6 and ID7. Table 7 shows the Mode 2 LUN priority. Figure 91 on page 116 shows the web map assigned sequence in a loop connection, Figure 92 on page 117 shows the web map assigned sequence in string a connection.

Mode 2 - LUN Priority

Table 7 (Page 1 of 2).

ID	LUNs
ID0	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID1	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID2	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID3	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID4	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID5	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID8	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7

Table 7 (Page 2 of 2).	
ID	**LUNs**
ID9	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID10	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID11	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID12	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID13	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID14	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7
ID15	LUN0 LUN1 LUN2 LUN3 LUN4 LUN5 LUN6 LUN7

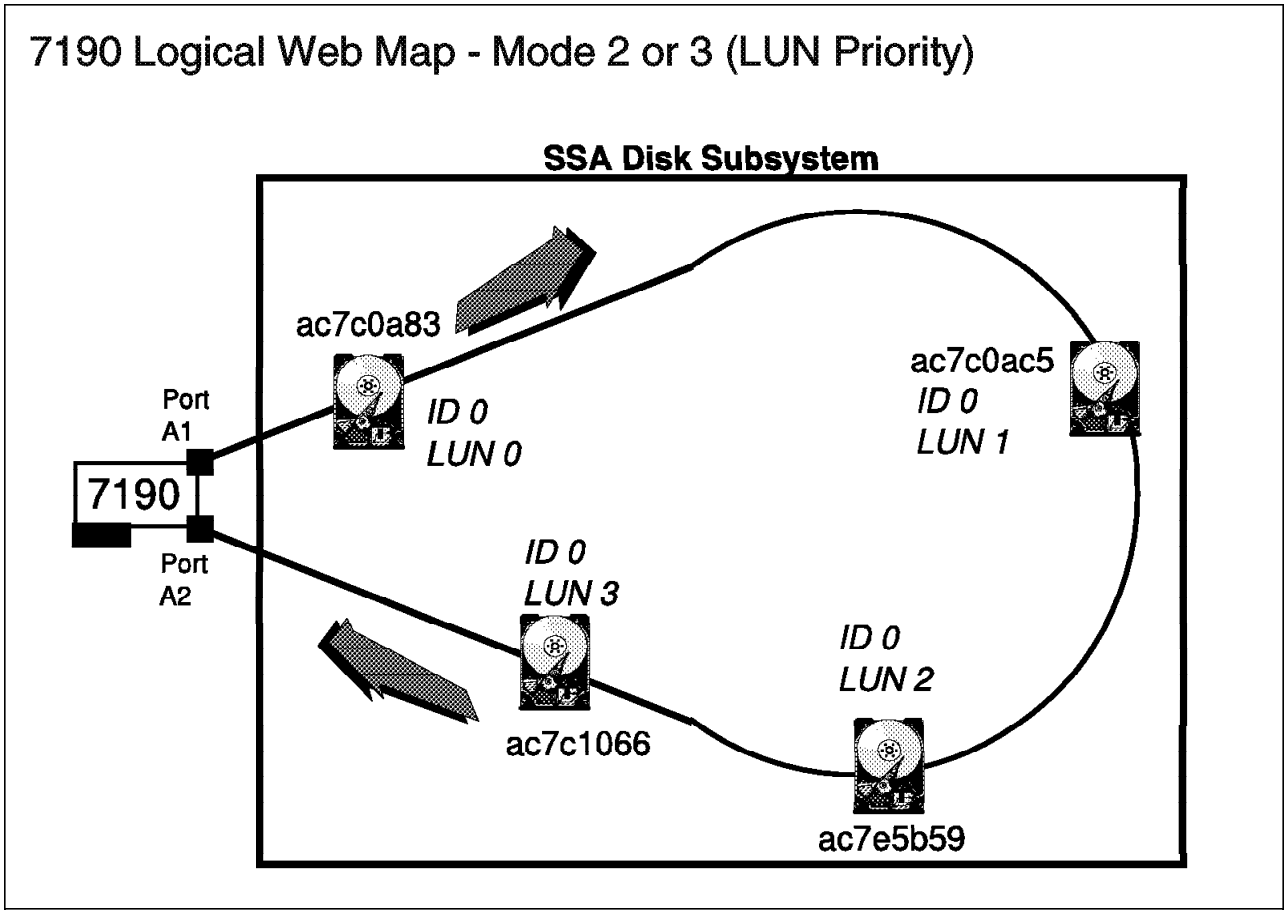

Figure 91. IBM 7190 Model 100 Mode 2 Loop Connection

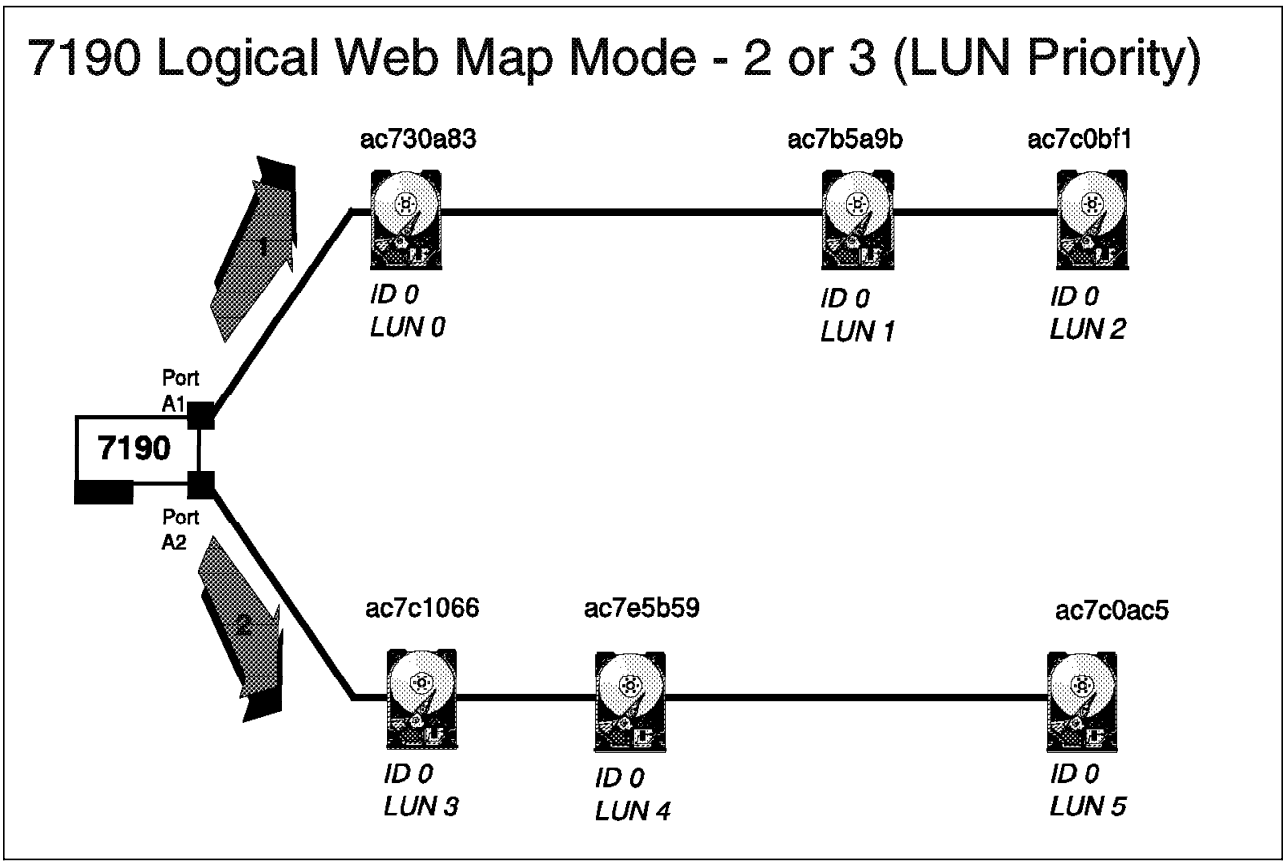

Figure 92. IBM 7190 Model 100 Mode 2 String Connection

Mode 3 is defined as LUN priority mapping, Fast/Wide SCSI, with device spin up. Mode 3 is the default in the 7190. The 7190 Model 100 uses LUN priority to map SSA drives to SCSI ID/LUN pairs. This method assigns SCSI ID/LUN pairs in the same way as Mode 2. The additional feature is issuing drive spin up commands to all devices on the SSA.

4.2.2.4 Software Installation

In this subsection we describe the process of copying files to disk and installing the 7190 Model 100 service functions application in Solaris.

Create 7190 Model 100 Partition on Solaris: The 7190 Model 100 service functions application requires a disk partition of at least 1 MB for its exclusive use. The partition must be created on the SSA devices attached to the 7190.

Installing Services Function on Sun: Installing the 7190 Model 100 service functions application requires root privileges. Because the utility is restricted an only system administrators. Follow these steps to install the 7190 Model 100 service functions:

1. If you are installing from a CD-ROM, follow these steps:
   a. Load the CD-ROM into the host system. Wait 10 seconds while the system carries out the device mount
   b. If the File Manager is running, it opens a window that displays the mount point and the contents of the CD-ROM. If the File Manager is not running, you can use the df command to display the mount point.
   c. Change to the mount point directory. For example,

cd /cdrom/cdrom0

d. When you execute the install program, the install is in the default directory, /ibm7190. If you want to install another directory, supply the directory name as a parameter to the install command. To execute the install program, issue this command:

./install

e. If the File Manager is running, select **Eject Disk** to unmount and remove the CD-ROM.

If the File Manager is not running, at the command line enter the following command to unmount the CD-ROM:

eject cdrom

and remove the CD-ROM from the system

2. If you are installing from diskette, follow these steps:

a. Place the 7190 Model 100 service functions diskette in the diskette drive and enter this command:

volcheck

b. Create a directory on the host system drive for the 7190 Model 100 service functions application. For example,

mkdir /ibm7190 model 100

c. Change to the directory you created. For example,

cd /ibm7190 model 100

d. Copy the tar files from the 7190 model 100 service functions application diskette with the following command:

tar -xvf/vol/dev/rdiskette0/unlabeled

where /vol/dev/rdiskette0/unlabeled is the special name assigned to the diskette drive by the volume manager.

e. When you execute the install program, the install is in the default directory, /ibm7190. If you want to install another directory, supply the directory name as a parameter to the install command. To execute the install program, issue this command:

./install

f. Remove the diskette with this command:

eject floppy

Before reading the sections that follow, you might want look at these Sun manuals for reference:

- *System and Network Administrator Guide*
- *Course material SM-101 Solaris 2.x Essentials for System Maintainers*

Creating 7190 Model 100 Configuration File on Solaris: For the 7190 Model 100 service functions application you can use a partition, but it must be specified in the 7190.cfg file located in the same directory as the install. With the system editor, create a file called 7190.cfg with the following structure:

SSA1 /dev/rdsk/cntndnsn

where SSA1 is the user-defined name, which can be up to 15 characters, and /dev/rdsk/cntndnsn is the device special file name for the primary 7190 Model 100 partition.

Partitioning the Disk Drives on Sun.: In this subsection we describe the process of partitioning and labeling each of the disks mapped by the 7190. You must use the SUN format utility. The format utility is a disk-maintenance tool used to partition disks as well as perform other disk-maintenance task. Here is the procedure:

1. Log onto the host system as root and enter the **format** command.

   When the format utility is invoked, it probes the SCSI bus to find the disk. All disk drives attached to the Sun system should be displayed, including the disk drives mapped by the 7190. Figure 93 on page 120 shows the output. If any of the expected SCSI ID/LUN pairs is not listed, refers to Chapter 5, “Maintenance” on page 125.

```
root> format

Searching for disk...done

AVAILABLE DISK SELECTIONS:

0.    cot3d0  <SUN0207 cyl 1254 alt 2 hd 9 sec 36>

       /sbus@1, f000000/esp@0, 800000/sd@3,0

1.    cot1d0  <SUN0669 cyl 1614 alt 2 hd 15 sec 54>

       /sbus@1, f000000/esp@0, 800000/sd@1,0

2.    cot2d0  <SUN0424 cyl 1151 alt 2 hd 9 sec 80>

       /sbus@1, f000000/esp@0, 800000/sd@2,0

Specify disk (enter its number):
```

Figure 93. Format Comand Output

2. Select the disk option for the SCSI ID/LUN with which you want to work.

   If the format utility prompts you with ″Disk not labeled. Label it now?″ enter an **n**, because the disk must be partitioned first

3. From the format menu (see Figure 94 on page 121), select the **partition** option and use the partition menu to set up partition sizes. (See Figure 95 on page 122). Set the **backup** partition to the total available cylinders. Reducing the size of partition 2 reduces the available disk space. Set the remaining partitions to utilize the disk space according to your requirements.

```
Selecting c0t3d0


FORMAT MENU:
disk - select a disk
type - select (define) a disk type
partition-select (define) a partition table
current - describe the current disk
format - format and analyze the disk
repair - repair a defective sector
analyze - surface analysis
defect - defect list management
backup - search for backup labels
verify - read and display labels
save - save new disk/ partition definitions
inquiry - show vendor, product and revision
volname - set 8-character volume name
quit
format> partition
```

Figure 94.

```
Selecting c0t3d0


PARTITION MENU:
0               - change '0' partition
1               - change '1' partition
2               - change '2' partition
3               - change '3' partition
4               - change '4' partition
5               - change '5' partition
6               - change '6' partition
7               - change '7' partition
select        - select a predefined table
modify      - Modify a predefined partition table
name         - name the current table
print          - display the current table
label          - write partition map and label to disk


partition> modify
```

Figure 95. Partition Menu

4. Write the partition information to the disk with the **label** option in the format menu. See Figure 95
5. Exit the partition menu and repeat the process for all disks mapped by the 7190.
6. When all disks have been partitioned and labeled, exit the format utility.

Creating and Mounting File Systems on Solaris.: In this section we describe the process of creating file systems by using the SUN **newfs** command. The **newfs** command allows you to initialize file system structures. Here is the procedure:

1. Use **newfs** command to create file systems for all partitions created on the SCSI ID/LUN pairs. For example:

   newfs /dev/rdsk/cnt1d0sn

   where sn is the partition created on SCSI ID 1, LUN 0.

Figure 96 on page 123 shows the output.

```
newfs: construct a new filesystem /dev/rdsk/cnt1d0sn: (y/n)? y

/dev/rdsk/cnt1d0sn: 28188 sectors in 87 cylinders of 9 tracks,
36 sectors

15.5MB in 3 cyl groups (16 c/g, 5.90MB/g, 2688 i/g)

super-block backups (for fsck -F ufs -o b= #) at:

32,  11632, 23232,
```

Figure 96. Newfs Command

2. The mount point directory must be created before any file system can be mounted. Create mount points for the partitions with the **mkdir** command. For example,

   mkdir /ssa1106

3. Mount the file systems to the directories created using the **mount** command. For example,

   mount /dev/rdsk/cnt1d0sn /ssa1106

   To verify that the mount was successful, use the **df -ak** command or the **mount** command. Figure 97 shows the output from the **mount** command.

```
root> mount

/          on  /dev/c0t3d0s0    read/write on Oct 28 18:33:50
1997

/usr     on  /dev/c0t3d0s6    read/write on Oct 28 18:33:50  1997

/home on  /dev/c0t3d0s7    read/write on Oct 28 18:33:50  1997

/ssa1106    on  /dev/cnt1d0sn    read/write on Oct 28 18:33:50
1997
```

Figure 97.

The operating system references the /etc/vfstab file (virtual file system table) at boot time to determine which file systems exist. You must edit the **/etc/vfstab** file to automatically mount the file systems when the system is booted. If you do not make a **vfstab** entry, it is necessary to unmount and

mount the file systems every time the system is shut down. Here is an example of the line that you would to **/etc/vfstab** to set up the sample file system to automatically mount at system startup:

/dev/rdsk/cnt1d0sn /dev/rdsk/cnt1d0sn /ssa1106 ufs 3 yes

Troubleshooting:: If you have trouble with the installation process, refers to Chapter 5, “Maintenance” on page 125

Utilities Description: Use the SSA Service Functions command, **SSA_SERV**, to perform maintenance, display the status of the SSA system, and diagnose problems with SSA disks, SSA adapters, and SSA system-level configurations.

The SSA_SERV utility uses a standard UNIX command line syntax, so the program is invoked from the shell prompt. To use SSA_SERV, the operator must be logged on as root.

To begin working with SSA_SERV, perform the next steps:

1. Change the directory that contains the SSA_SERV program and the 7190.cfg file. Refer to 4.2.2.4, “ Software Installation” on page 117
2. Issue this command:
   /SSA_SERV

   or

   /SSA_SERV name_7190

   or

   /SSA_SERV -dname_7190 model 100 -stime_sec

   where the **-d** option is used to specify the name of the 7190 Model 100 as defined in 7190.cfg during installation and **-s** is used to specify the sample time in seconds. The sample time allows collecting errors from the 7190. The errors are logged into the SSA service functions error log. If you set the sample time to 0, the error logging function is disabled.

   The utility displays the menu shown in Figure 98 on page 126.

3. From the SSA_SERV menu, select **t.** (*”*Show SSA topology*”*). The disk drive are shown in the sequence they are physically attached to the SSA loop or string.

   Detailed information can be found in:

   - *IBM Manual 7190 Model 100 Host to SSA Loop Attachment Model 100, Installation and User’s Guide, GC26-7195-0.*
   - Chapter 5, “Maintenance” on page 125..

Chapter 5. Maintenance

This chapter addresses topics such as troubleshooting, upgrading and replacing hardware, upgrading microcode, and other maintenance requirements.

5.1 7190 Model 100 operations

This section describes all the procedures relating to the 7190. It is divided into six subsections:

- **Monitoring the SSA network** deals with the functions that provide information about the network without changing the characteristics of the network.
- **Maintaining the SSA web** deals with functions that provide information about the SSA network and functions that change the characteristics of the SSA network.
- **Adding drives** describes the procedures to add SSA disk drives to the various SSA configurations.
- **Replacing drives** describes the procedures to replace drives within the various SSA configurations.
- **Loading firmware** describes the microcode download procedures for the SSA network components.
- **Troubleshooting** describes the functions available for system troubleshooting and provides a problem determination procedure.

All the procedures in this section are related to the functions available through the **SSA Service Functions** menu. Start the SSA Service Functions as follows:

1. From the Solaris command shell, change to the directory that contains the SSA_SERV program and the 7190.cfg file. (The program and cfg file are installed during the installation procedure.)
2. Type the command to start the service functions : **./SSA_SERV -dnnn** (replace **nnn**with the name of the adapter as in the configuration file)

   or:

   ./SSA_SERV -dnnn -sttt (replace nnn with the name of the 7190 Model 100 and replace ttt with the sample time required)

   The **-d** option is used to specify the name of the 7190 Model 100 as defined in the 7190.cfg file. The **-s** is used to specify the sampling time. The sampling time is the interval between each operation to collect errors from the 7190 Model 100, which are logged into the SSA service functions error log. A time setting of zero disables the log.

   Note

   If access to all 7190 Model 100 partitions is lost, the following following message is displayed:

   ```
   >>Lost access to all 7190 model 100 partitions
   ```

   The SSA Service functions menu is displayed in Figure 98 on page 126

```
1. Download Microcode
2. Display VPD
3. Identify a Disk Drive Module
4. Set/Reset Service mode
5. Disk Drive Diagnostics
6. 7190 model 100 Diagnostics
7. Error log Analysis
8. Clear Check mode
9: Activities Monitor
t. Show SSA Topology
m. Force SSA master
r. Force Web reset
w. Write Cache Option
s. Synchronous Write
q. Quit
?. Menu
>>
```

Figure 98. Menu of SSA_SERV Options

Any of these options can be selected by typing the number or letter that identifies it.

5.1.1 Monitoring the SSA Network

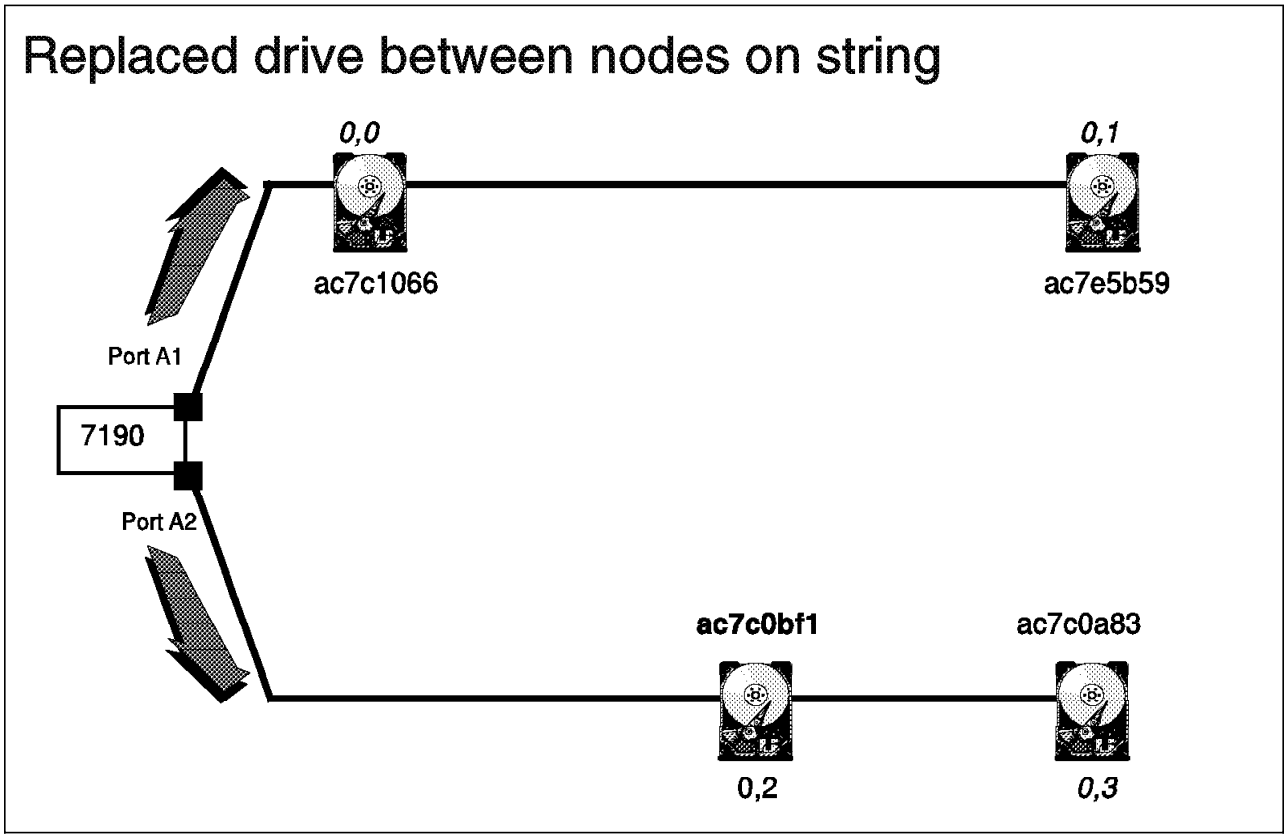

Figure 99. Base Configuration Example

We suggest that you file the output of the commands in this section. The information can prove valuable when troubleshooting your SSA network.

5.1.1.1 Adapter Operations

This section describes the functions that you can use to gather information about the adapter.

Display VPD: This function displays the vital product data (VPD) of the selected adapter. From the SSA Service Function menu, select **2. Display VPD** This selection will give you the option of directing the command to either the 7190 Model 100 or one of the SSA drives attached:

```
>> 2

>> Input selection (s for 7190 model 100, TxDy for drive ID x and
LUN y): s
```

Choose **s** to select the 7190. This will display information similar to the following:

```
>>
>>Vital Product Data of selected device:
>>Vendor ID - IBM
>>Product type - 7190 model 100
>>Model Number - I
>>Microcode Revision - 4.1
>>Unit Serial Number - 12345678
>>SSA Unique ID - 0000006022000000
>>
```

Show SSA Topology: This function displays a logical representation of the association of SCSI ID/LUNs with SSA UIDs.

From the SSA Service Functions menu, select **t. Show SSA Topology**.

The disk drives are listed in the same sequence as they are physically attached to the SSA Loop or SSA string:

```
nnn7190.22000000

SSA.PORT.A1=
=I.T0D0.AC7C1066.OK=
=N.T0D1.AC7E5B59.OK=
=N.T0D2.AC7C0AC5.OK=
=N.T0D3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2

INACTIVE:
<T0D5.AC7B5A9B.OK>
```

The columns of information on the screen can be categorized:

- **nnn.7190.uuuuuuuu.SM** (.7190.220010BA.SM= in the screen shot)

  The characters **nnn** represent the name of the 7190 Model 100 which can be up to 15 characters long. The name is followed by the product number **7190**

and then the last eight digits **uuuuuuuu** of the SSA UID for the attachment hardware. The status code **SM** is last.

- hp2.SSA.PORT.A1

  This indicates whether the port is attached to a drive. The **=** indicates the port connector A1 is connected to the following drive. If there is no drive connected to the port, the line reads **SSA.PORT.A1>**

- SSA.PORT.A2

  The SSA port connector A2 on the 7190 Model 100

- ?mTxDy.uuuuuuuu.ss&

  We will use **=N.T0D1.AC7E5B59.OK=** from the Show SSA Topology output as our example:

  – **?** Port connection (= in our example): **=** means the port is connected to the drive **<** means the port is the end of the string connected to port 2 of the 7190.

  – **m** mode: **S** The disk drive module is in Service mode. **I** The disk drive module is in identify mode. **N** The disk drive module is in normal operating mode.

  – **Tx** Target ID (0 in our example)

  – hp2.Dy Target LUN (1 in our example)

  – **uuuuuuuu** The last eight characters of the SSA UID (AC7E5B59 in our example)

  – ss The status of the SSA Node (OK in our example): **OK** The disk drive module is in normal state. **PC** The disk drive module is detecting a loss of redundant power or cooling. **CK** The SSA service functions have logged an error for this device (The device may or may not be operational.)

  SM The currently selected 7190 Model 100 is the master on the SSA web.

  – **&** Port connection (= in our example): **=** The port is connected to the drive or port A2 specified in the previous line.

  This port is the end of the string connected to port 1 of the 7190.

- **INACTIVE**(AC7B5A9B in our example): Shows drives that are not currently configured in the loop or string because they have been physically removed or because they have been placed in service mode.

Activities Monitor: This function displays the current activities of the SCSI port and SSA port1 and port 2 on the 7190 Model 100 attached to an SSA network.

Select **Activities monitor** from the main menu:

```
SCSI :        RMBPS= XX.XX    WMBPS= XX.XX IOPS= XX.XX
SSA.PORT.A1 : RMBPS= XX.XX    WMBPS= XX.XX
SSA.PORT.A2 : RMBPS= XX.XX    WMBPS= XX.XX

Hit return to terminate monitor.
```

The columns of information can be categorized as follows.

- **SSA.PORT.A1** The SSA port connector A1 on the 7190 Model 100
- **SSA.PORT.A2** The SSA port connector A2 on the 7190 Model 100
- **RMBPS** The read data transfer rate in megabytes per second. The displayed value is the average of the last 16 samples.
- hp2.WMBPSWrite data transfer rate in megabytes per second. The displayed value is the average of the last 16 samples.
- hp2.IOPS.ehp2. Number of I/O operations per second. The displayed value is the average of the last 16 samples.

5.1.1.2 Disk Drive Operations

This section describes the functions available for you to monitor disk activities.

Display VPD: If you had chosen **t0d0** (for a disk drive) for diagram Figure 111 on page 142, the following information would have been displayed:

```
>>
>>Vital Product Data of selected device:
>>Vendor ID - IBM
>>Product type - DCHC
>>Model Number - O9B
>>Microcode Revision - 7777
>>Unit Serial Number - 13062975
>>SSA Unique ID - 0000000AC7E5B59
>>
```

You can return to the SSA Service Functions main menu at any time by typing **?**.

Identify a disk drive module: This function is used to identify the location of any number of drives attached to the 7190. Identify causes the check LED on the drive as well as its enclosure (This applies to the 7133 and 7131 subsystems) without disrupting normal operation of the drive. For the purpose of our example, we assume that the 7190 Model 100 is in default operation Mode 3.

Select **3.** (″Identify a Disk Drive Module″) from the SSA Service Functions menu. his will ask you to select a target:

```
>>3

>>Input selection (TxDy, for drive ID x and LUN y):
t0d0
>>
```

Once you have made your selection, the Show SSA Topology display appears:

```
*** Show SSA TOPOLOGY ***

nnn7190.22000000

SSA.PORT.A1=
=I.TOD0.AC7C1066.OK=
=N.TOD1.AC7E5B59.OK=
=N.TOD2.AC7C0AC5.OK=
=N.TOD3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2

>>
```

The check LED flashes on the selected drive (see **ID 0 LUN 0** in Figure 111 on page 142) and the Show SSA topology command displays an **I** next to the drive.

To **end identify**, select identify again and target the same drive.

5.1.2 Maintaining the SSA web

5.1.2.1 Display the Logical Web Map

5.1.2.2 Force SSA Master

This function forces the selected 7190 Model 100 to become master of the SSA network and resets all current service mode drives. Only the master device has service mode setting ability.

Choose **m.** (″Force SSA Master″) from the main menu:

```
>>m

>>Are you sure that you want to force the local 7190 model 100 to be the SSA
Master ? (type Y or N):y

>>done.
```

5.1.2.3 Force Web Reset

This function forces a web reset to all the devices on the web. Web reset can be ordered only by the 7190 Master:

```
>>

>>Are you sure that you want to do an SSA Web Reset? (type Y or
N):y

Start Web Reset

>>Done.
```

5.1.3 Adding Drives to a 7190 Model 100 Network

This section deals with expansion of 7190 Model 100 networks, drive removal, and drive replacement. Be sure to follow exact procedures to avoid data loss. Always have a full backup available before altering the structure of the SSA network.

When a new drive is installed into the SSA network, it is assigned the next available SCSI ID/LUN combination. When a drive is replaced, its SCSI ID/LUN combination is transferred to the replacement drive. .

5.1.3.1 Adding Drives to a 7190 Model 100 loop

This procedure requires you to make a single break in the SSA loop. Be sure that there is only one break in the loop while the system is running as multiple breaks can isolate drives from the 7190 Model 100, resulting in possible data loss.

Multiple drives can be added simultaneously at the same break. If you need to add drives to different positions on the SSA network, repeat this procedure for each drive.

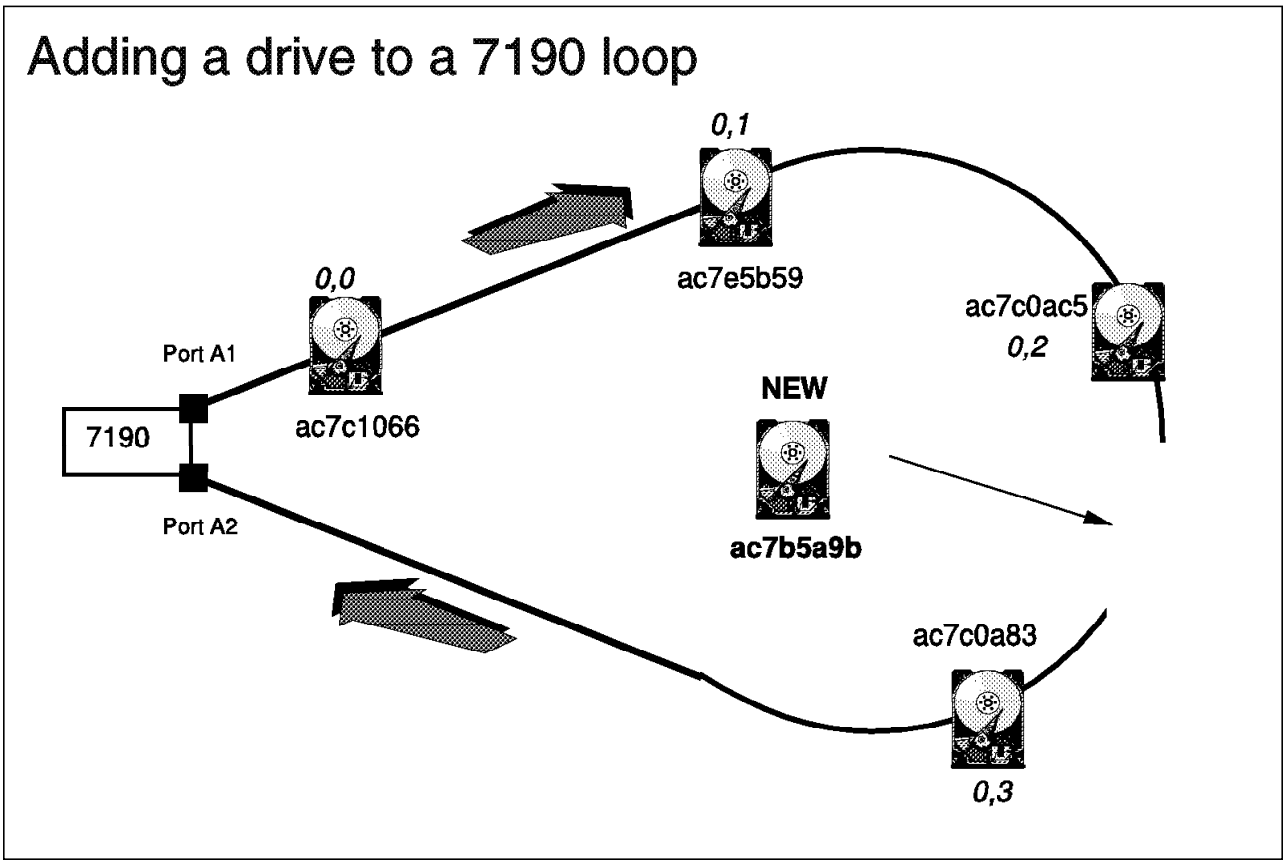

Figure 100. Adding a Single Drive to a 7190 Model 100 SSA Loop

Use the following procedure:

1. Enter the **SSA_SERV** utility and view the SSA web topology, verifying that the loop is complete and no drives are in the inactive list.
2. Identify where to split the loop. In our example (see Figure 100), we have chosen to insert the drive with UID **ac7e5b59** between ***0,2*** and ***0,3***.

3. Split the loop into two strings and attach the new drive.
4. Your altered configuration should look similar to that in Figure 101.

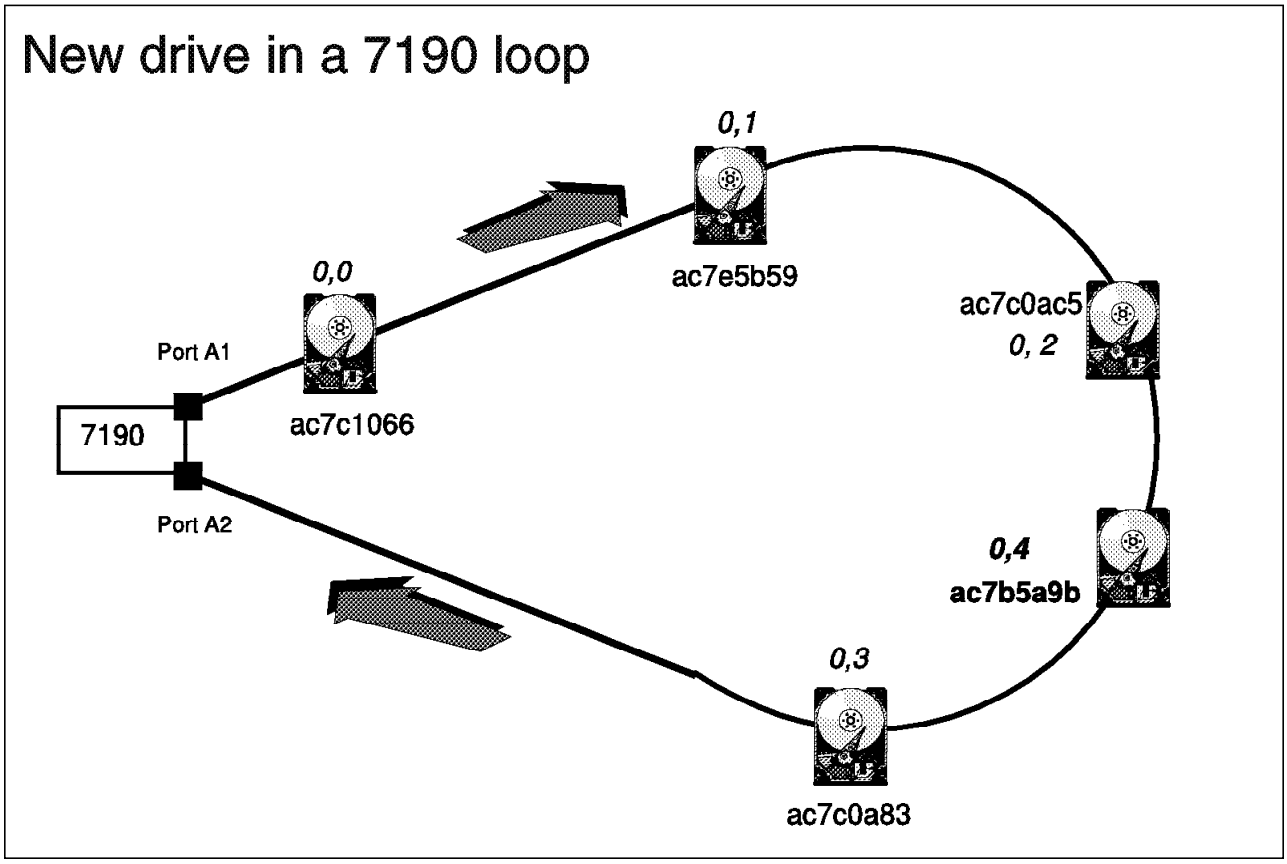

Figure 101. The New Configuration

5. Verify that the 7190 Model 100 recognizes the new drive by displaying a web topology:

```
nnn7190.22000000
SSA.PORT.A1=
=N.T0D0.AC7C1066.OK=
=N.T0D1.AC7E5B59.OK=
=N.T0D2.AC7C0AC5.OK=
=N.T0D3.AC7C0A83.OK=
=N.T0d4.AC7B5A9B.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2
```

Compare the output to Figure 101.

6. Reboot Solaris to access the new drive:

   root> reboot -- -r (from the prompt)

   or:

   boot -r (from the openboot prompt)

7. The new drive can now be partitioned and made ready for use.

Note

If you run a procedure that causes a **Node Map Reset**, SCSI ID/LUN combinations are reassigned according to the mapping rule, thus affecting applications that rely on specific SCSI ID/LUN maps. We strongly recommend that you add a new drive as the node **closest to port 2** of the 7190 Model 100 to avoid remapping during a reset.

We recommend that you add all new drives to the end of an SSA string to avoid interrupting access to downstream nodes and to avoid potential problems caused by remapping during a reset.

5.1.3.2 Adding Drives to a 7190 Model 100 String

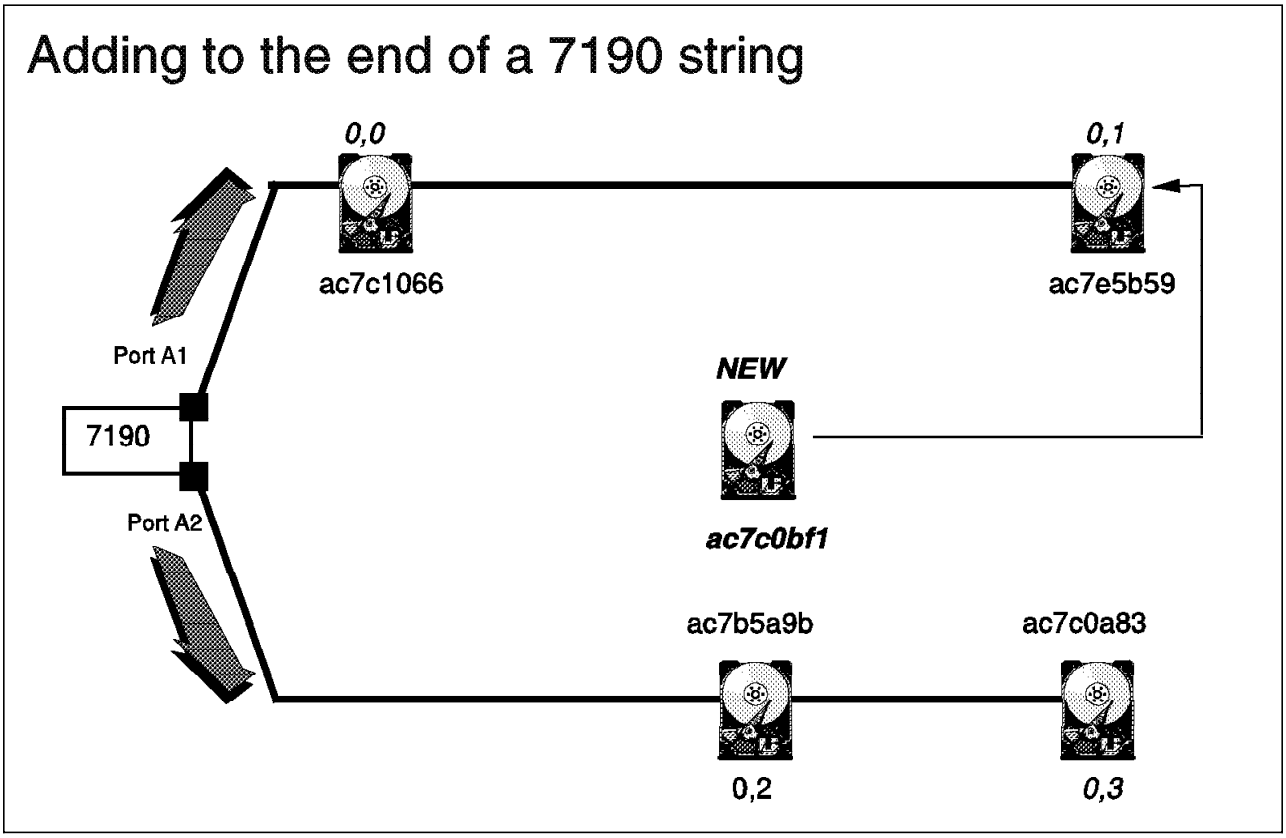

Figure 102. Adding SSA Drives to the End of a 7190 Model 100 String

Use the following procedure:

1. Display the web topology to verify that all the drives on the string are active.
2. Identify the string to which you will add the new drive. (In our example, we will add the drive with UID **ac7c0bf1** to the end of string 1 which originates from port 1 of the 7190.)
3. Attach the new drive.

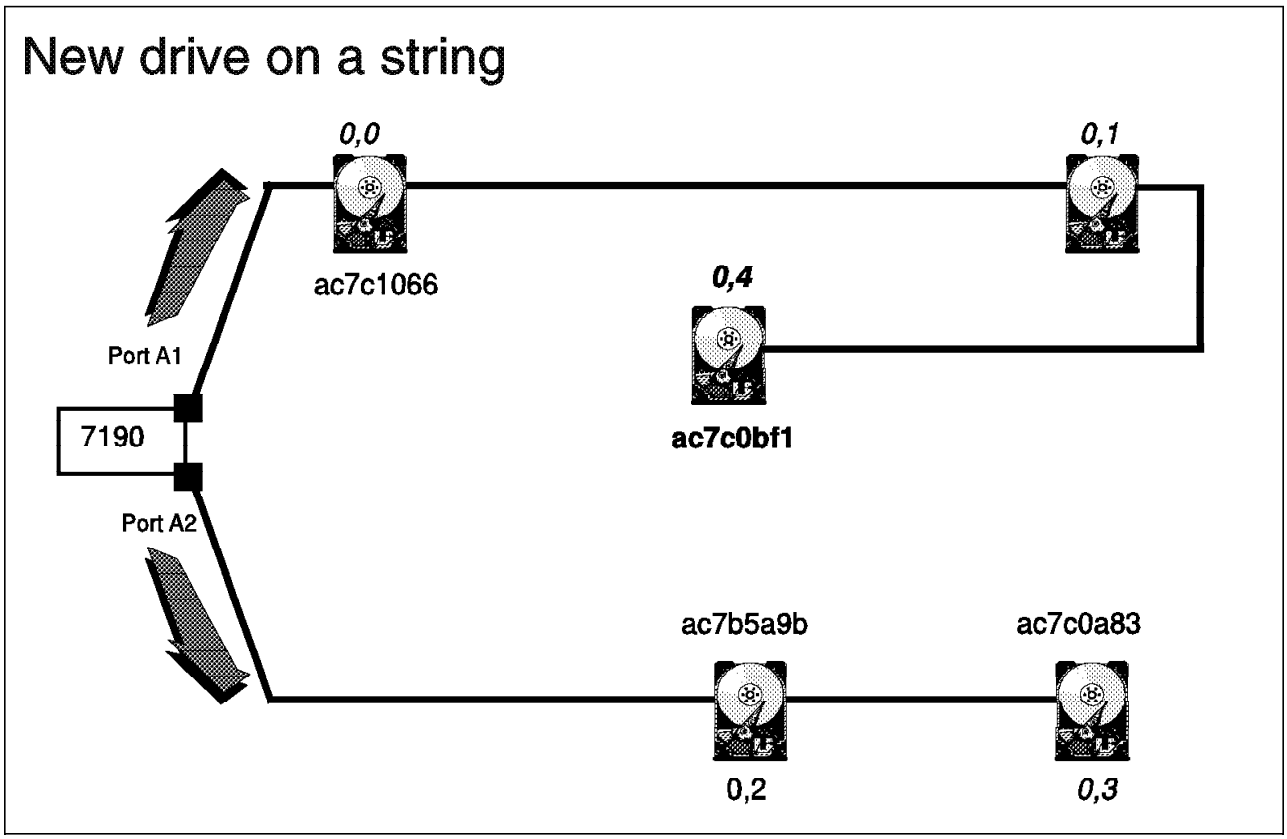

Figure 103. The Newly Added Drive at the End of String 1

4. Display the SSA web topology:

```
nnn7190.22000000
SSA.PORT.A1=
=N.T0D0.AC7C1066.OK=
=N.T0D1.AC7E5B59.OK=
=N.T0D4.AC7C0BF1.OK=
=N.T0D2.AC7B5A9B.OK=
=N.T0d3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2
```

   Compare the output with Figure 103.

5. Reboot Solaris to access the new drive:

   root> reboot -- -r (from the prompt)

   or:

   boot -r (from the openboot prompt)

6. The new drive can now be partitioned and made ready for use.

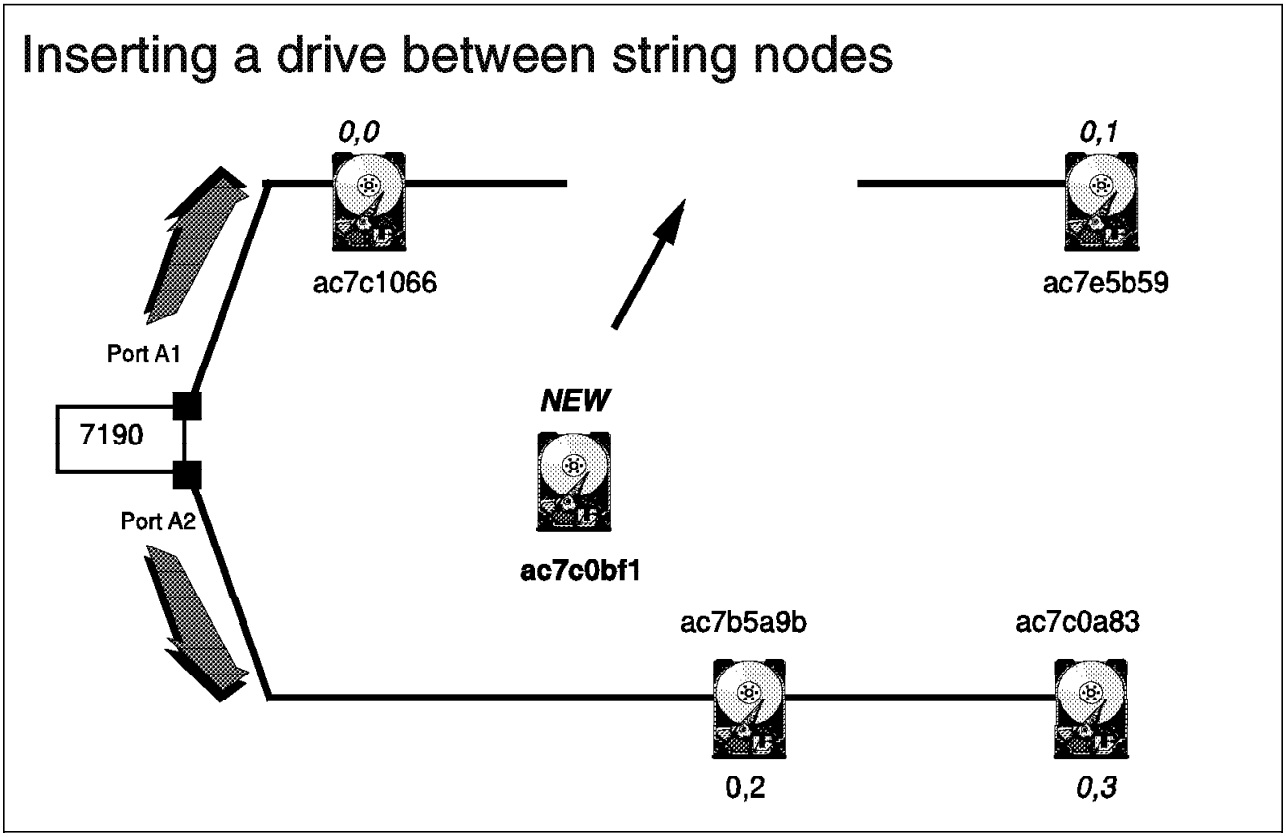

Figure 104. Inserting a Drive Between Nodes on a 7190 Model 100 String

5.1.3.3 Inserting Drives Between Existing Nodes on a String

If you have a specific requirement to insert drives between nodes, use this procedure:

1. Shut down Solaris and power down the system
2. Break the string in the chosen position. In our example, we have broken the string between ***0,0*** and ***0,1*** (Figure 104)
3. Add the new drive and power it up.
4. Power up the system.

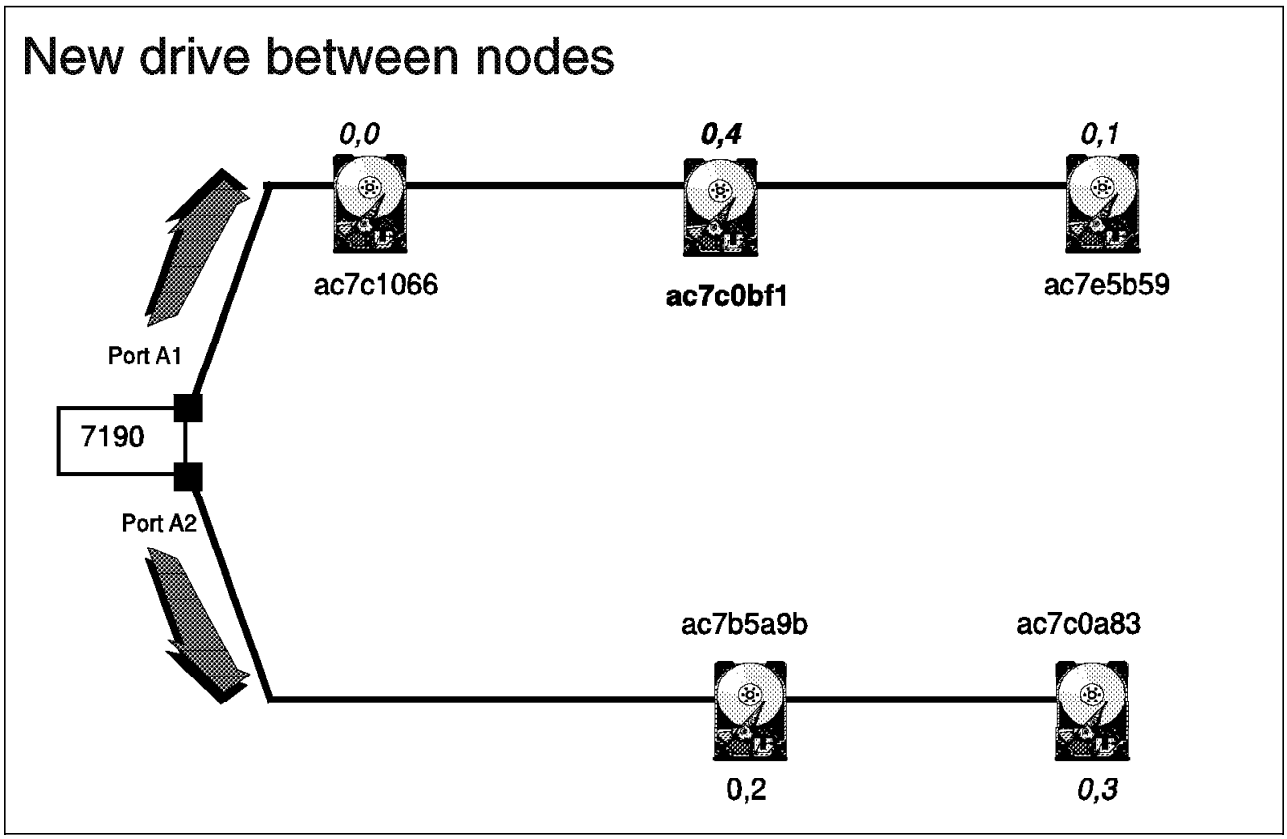

Figure 105. The Newly Added Drive Between the Nodes of String 1

5. Display the SSA web topology:

```
nnn7190.22000000
SSA.PORT.A1=
=N.T0D0.AC7C1066.OK=
=N.T0D4.AC7C0BF1.OK=
=N.T0D1.AC7E5B59.OK=
=N.T0D2.AC7B5A9B.OK=
=N.T0d3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2
```

   Compare the output to Figure 105.

6. Reboot Solaris to access the new drive:

   root> reboot -- -r (from the prompt)

   or:

   boot -r (from the openboot prompt)

7. The new drive can now be partitioned and made ready for use.

5.1.4 Replacing Devices on a 7190 Model 100 Network

Only one drive can be replaced at a time on an active system. The first access to a new device being powered up may cause the Solaris target drive to display error messages because the device has not yet spun up. A subsequent access should correct the error by spinning the drive up.

5.1.4.1 Replacing the 7190 Model 100

You may need to replace the 7190 Model 100 if you suspect it is faulty. If you have added new drives between existing drives since the last web reset, some of the drives on the 7190 Model 100 network will be remapped to new SCSI ID/LUN pairs. It is important to understand the logical mapping of your applications and user shares to physical SSA drives so that you can make the required alterations if a remap occurs.

The replacement procedure requires you to halt I/O to all devices attached to the 7190 model 100 to be replaced, and to shut down and power off the host system.

Once the host system is down, replace the 7190 model 100.

5.1.4.2 Replacing Drives in a 7190 Model 100 Loop

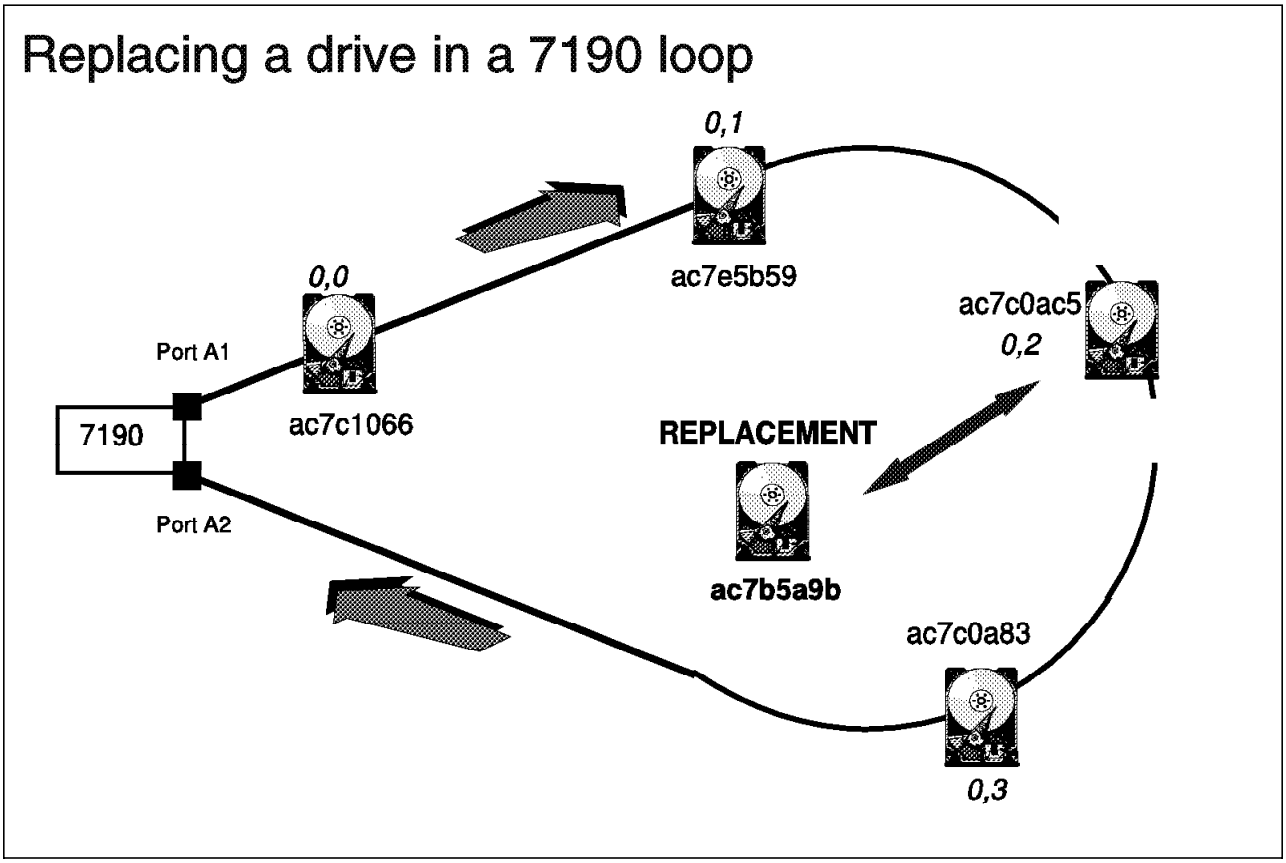

Figure 106. Replacing a Drive on an SSA Loop

Use the following procedure:

1. Take necessary steps to make sure that all I/O to the drive to be replaced is halted. The safest way to do this is to unmount the associated file systems.

2. Enter the SSA_SERV utility and view the SSA topology Figure 106) to verify that the loop is complete and no drives are in the inactive list. If the drive to be replaced is in the inactive list, you can skip Step 3.
3. Set the drive to **service mode**. This assures that:
   - Removing the drive will not isolate the other drives in the loop.
   - The drive motor is stopped.
   - The check LED on the drive is ignited.
   - The drive is automatically put into the inactive list.
4. Replace the drive making sure to close the loop.
5. Reset service mode on the drive in order to:
   - Turn off the check LED.
   - Start the drive motor.
   - Put the drive into the active list.

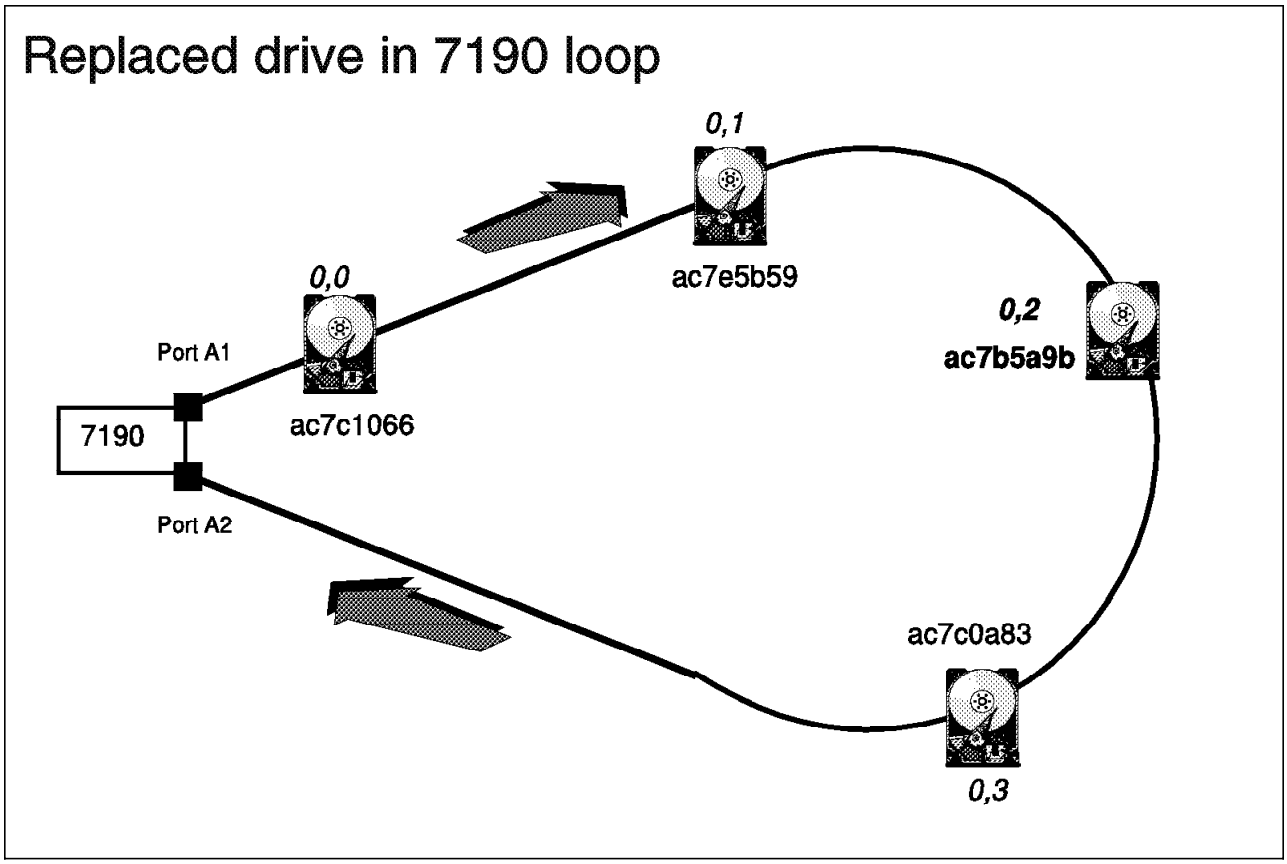

Figure 107. The Replaced Drive in the 7190 Model 100 Loop

6. Verify that the 7190 Model 100 recognizes the replacement drive by displaying the web topology:

```
nnn7190.22000000
SSA.PORT.A1=
=N.TOD0.AC7C1066.OK=
=N.TOD1.AC7E5B59.OK=
=N.TOD2.AC7B5A9B.OK=
=N.TOd3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2
```

Compare the output to Figure 107 on page 138

7. Reboot Solaris to access the new drive:

   root> reboot -- -r (from the prompt)

   or:

   boot -r (from the openboot prompt)

8. The new drive can now be partitioned and made ready for use.

5.1.4.3 Replacing Drives on a 7190 Model 100 String

There are separate procedures to replace a drive at the end of a string and to replace a drive between nodes on a string, the latter procedure requiring two reboots.

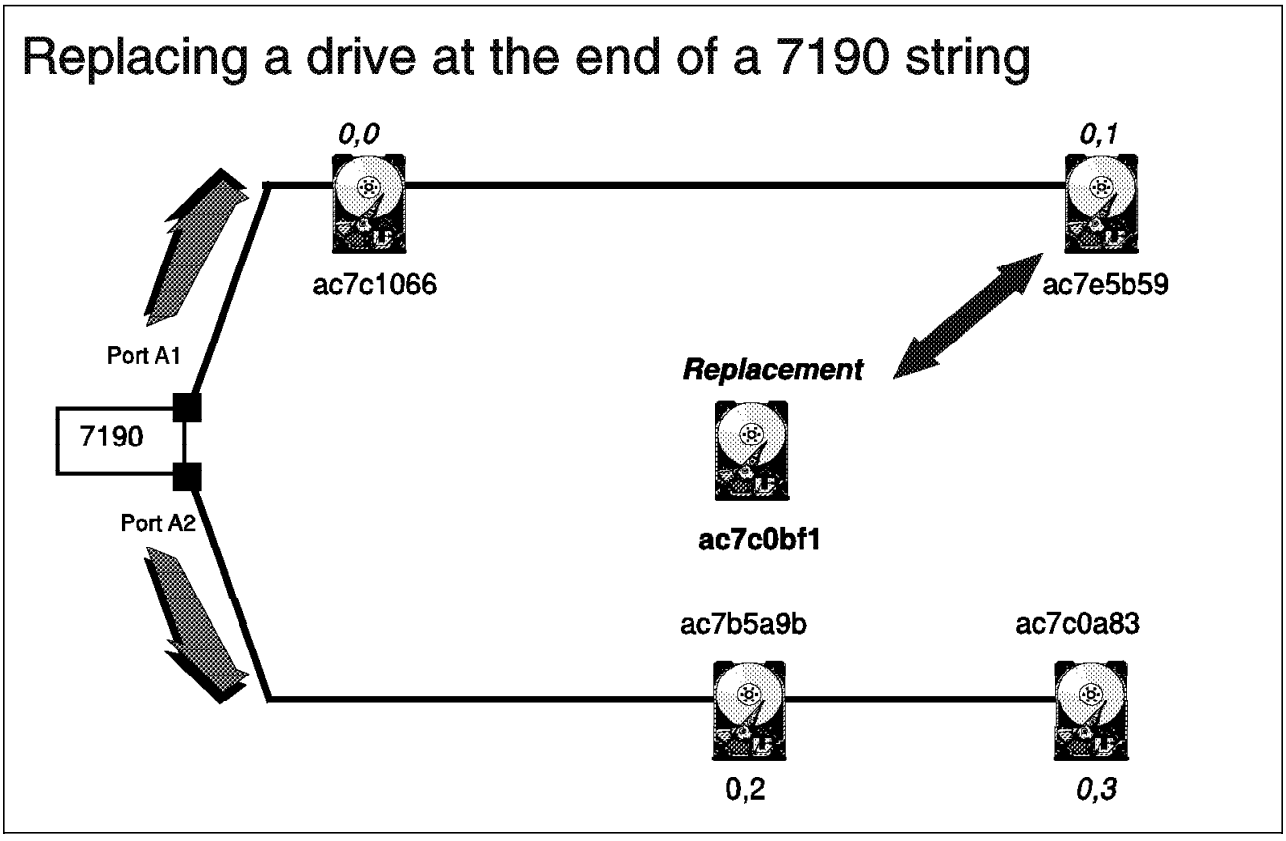

Figure 108. Replacing an SSA Drive at the End of a 7190 Model 100 String

Use the following procedure:

1. Take the necessary steps to halt all I/O to the drive to be replaced. The safest way to achieve this is to unmount the associated file systems.
2. Enter the SSA_SERV utility and view the web topology making sure that no drives on the string are in the inactive list. If the drive to be replaced is in the inactive list, skip Step 3.
3. Set the drive to **service mode**. This assures that:
   - Removing the drive will not isolate the other drives in the string.
   - The drive motor is stopped.
   - The check LED on the drive is ignited.
   - The drive is automatically put into the inactive list.
4. Replace the drive.
5. Reset service mode on the drive in order to:
   - Turn off the check LED.
   - Start the drive motor.
   - Put the drive into the active list.

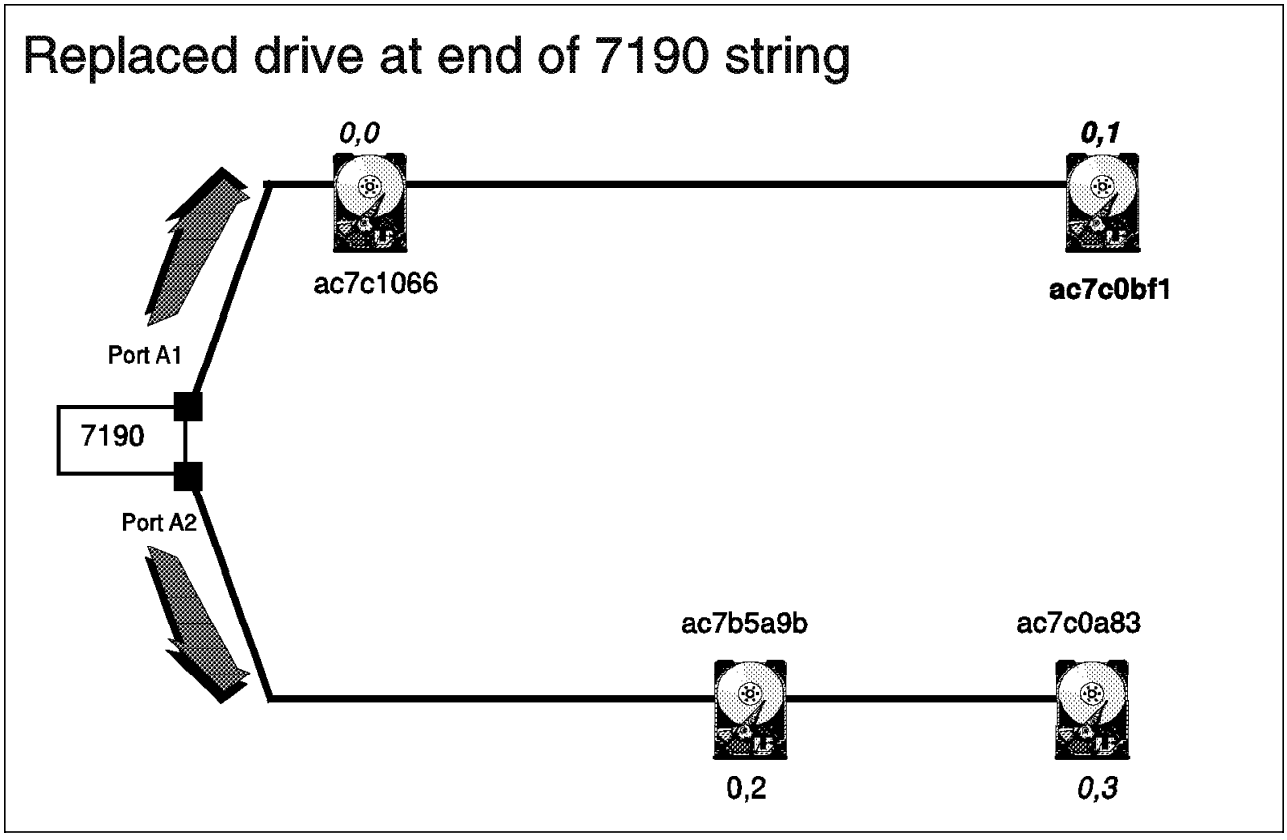

Figure 109. The Replaced Drive at the End of the 7190 Model 100 String

6. Verify that the 7190 Model 100 recognizes the replacement drive by displaying the web topology:

```
nnn7190.22000000
SSA.PORT.A1=
=N.TOD0.AC7C1066.OK=
=N.TOD1.AC7C0BF1.OK=
=N.TOD2.AC7B5A9B.OK=
=N.TOd3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2
```

Compare the output to Figure 107 on page 138

7. Reboot Solaris to access the new drive:

   root> reboot -- -r (from the prompt)

   or:

   boot -r (from the openboot prompt)

8. The new drive can now be partitioned and made ready for use.

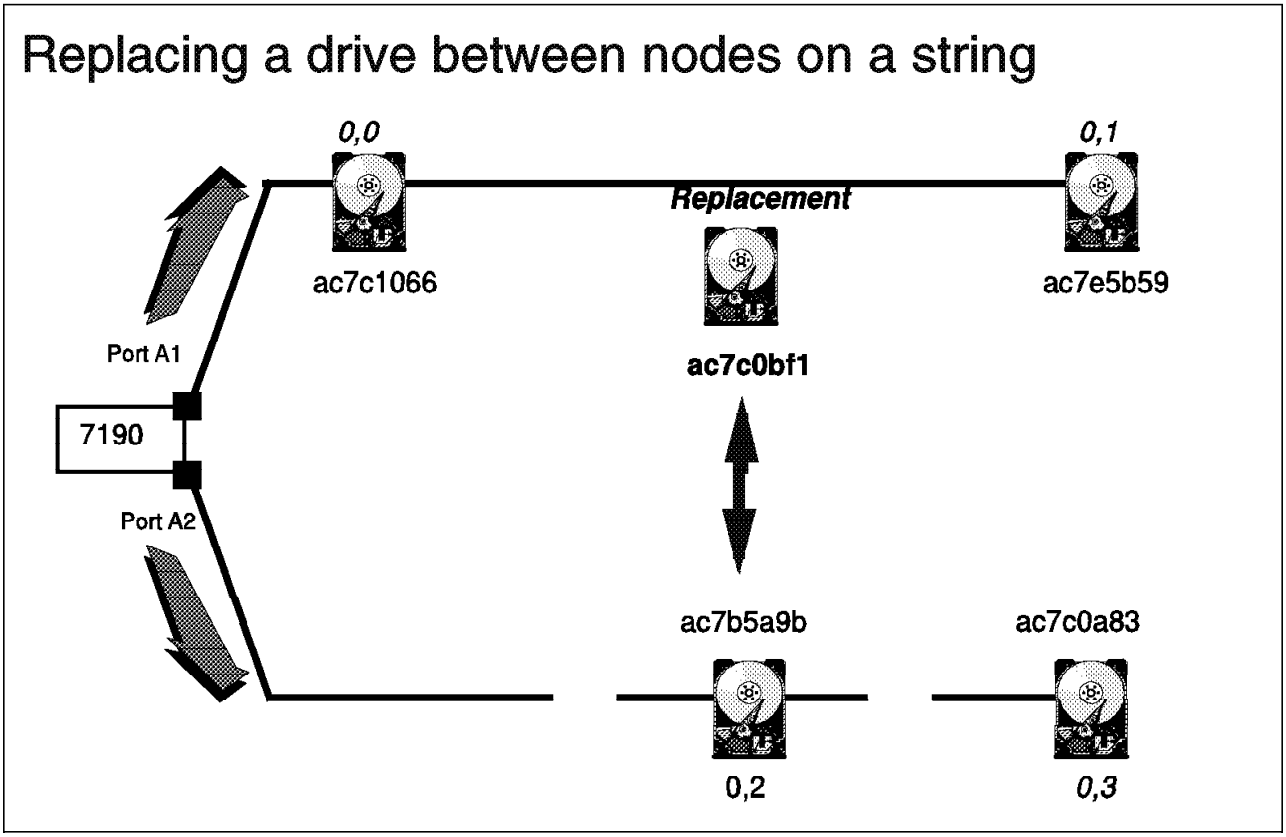

Figure 110. Replacing a Drive Between Nodes on a String

1. Shut down Solaris and power down the system.
2. Replace the drive.
3. Power up the system.

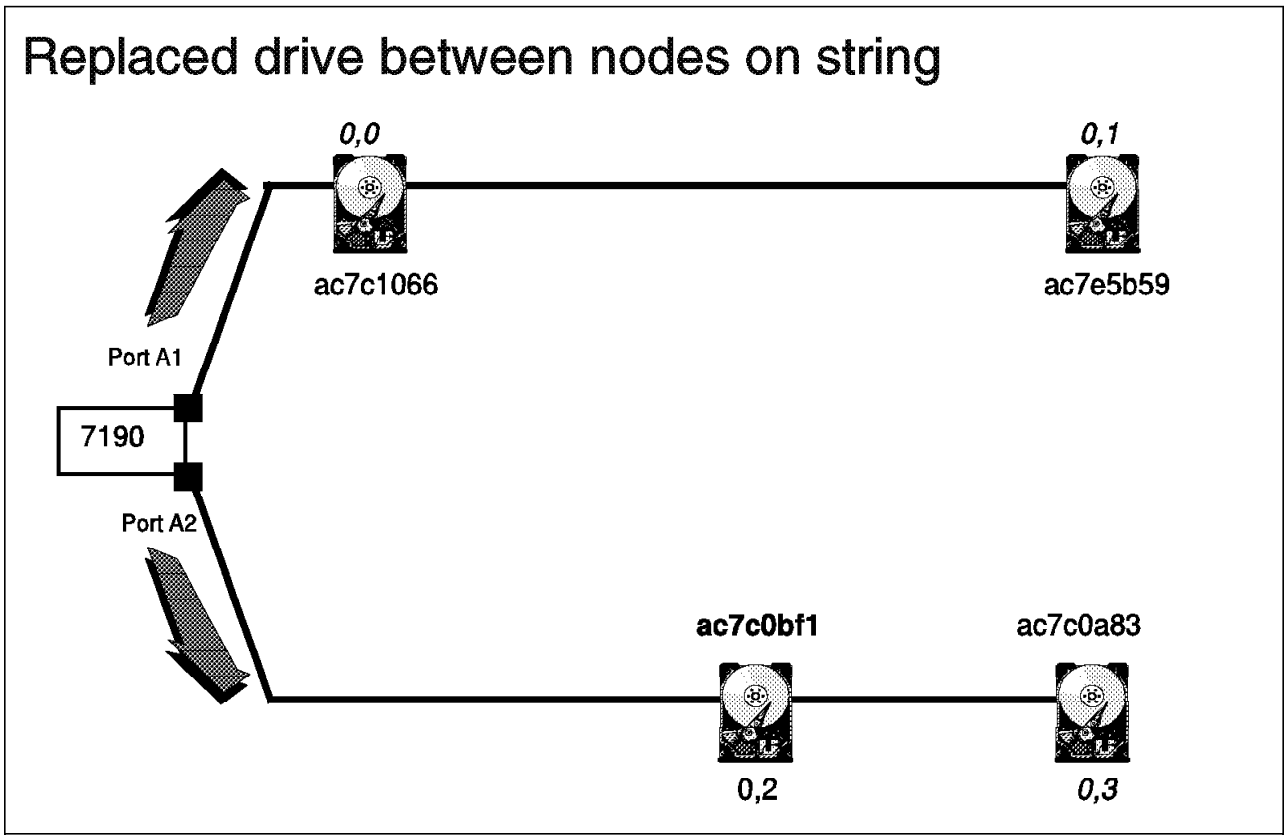

Figure 111. Replaced Drive Between Nodes on a 7190 Model 100 String

4. Display the SSA web topology:

```
nnn7190.22000000
SSA.PORT.A1=
=N.T0D0.AC7C1066.OK=
=N.T0D1.AC7E5B59.OK=
=N.T0D1.AC7C0BF1.OK=
=N.T0d3.AC7C0A83.OK=
=.7190.220010BA.SM=
=SSA.PORT.A2
```

   Compare the output to Figure 111.

5. Reboot Solaris to access the new drive:

   root> reboot -- -r (from the prompt)

   or:

   boot -r (from the openboot prompt)

6. The new drive can now be partitioned and made ready for use.

5.1.5 Loading Firmware

This section explains how to update firmware on the 7190 Model 100 or SSA drives attached to the 7190. We advise you to do firmware updates only on systems that are backed up by a UPS. Any power cycle during a firmware download to either the 7190 Model 100 or SSA drives can cause damage. Only the local 7190 can be downloaded in a multi-initiator environment. Make sure that the 7190 Model 100 is set to SSA master by viewing the web topology display.

Use the following procedure to update firmware on either the 7190 Model 100 or an SSA drive attached to the 7190:

1. Verify the exact path for the file that contains the firmware image.
2. Stop all SSA disk activity on the SSA network attached to the target 7190.
3. From the SSA Service Functions menu, select **download microcode**:

```
>>1
>>
>>Input selection (S for 7190 model 100, TxDy for drive id x and
lun y, æAÆ for all drives):t0d2
>>Download file name
e.g./usr/tmp/DFHC6969.DL):/usr/tmp/DFHC7171.DL
>>Option (N for normal, F for force):n
>>
```

4. Type your selection and file name. In the screen example shown here, we have chosen **t0d2** as the target and **/usr/tmp/DFHC7171.DL** is the file name of the firmware image, and *n* (normal) mode has been selected:
   - If normal mode is selected, only the devices that have a lower level of microcode will be targeted for the procedure. If **f (forced)** mode is selected, all the selected devices will be targeted for the procedure.
   - The file type must be **.DL**. A different file type will cause an error message suggesting the code is bad.
5. An automatic reset is issued after the download is completed. A successful download will cause the following to be displayed:

```
>>Download in process..............Done.
>>
```

6. The procedure is complete.

5.1.6 Troubleshooting

5.1.6.1 Disk Drive Diagnostics

Use this function to run diagnostics on drives which you suspect are faulty. Make sure that all I/O to the drive to be tested has been halted.

1. Select **5. Disk Drive Diagnostics** from the SSA Service Functions menu. This option will prompt you for the target drive:

```
>>Input selection (TxDy for drive ID x and LUN y):t0d4
```

2. Type your selection. In the example screen, disk drive t0d4 has been selected:

```
>>Input selection (TxDy for drive ID x and LUN y):t0d4
Are you sure that the selected device is not in use by any of the
attached systems? (type Y or N): y
```

   If you type *y*, the diagnostics run and progress messages are displayed:

```
>>Input selection (TxDy for drive ID x and LUN y):t0d4
Are you sure that the selected device is not in use by any of the
attached systems? (type Y or N):y
>>Test Start...Stop .Unit..Read Capacity..Read Block max..Done.
```

3. The test is complete

5.1.6.2 7190 Model 100 Diagnostics

Use this option to run diagnostics on the 7190 model 100 attachment. All I/O to the drives connected to the 7190 Model 100 must be halted.

1. Select **6. 7190 Model 100 100 Diagnostics** from the SSA Service Functions menu. The following prompt will be displayed:

```
>>Are you sure that all the drives on this 7190 Model 100
are not in use? Type Y or N:y
```

2. If you type *y*,the diagnostics run and progress messages are displayed:

```
>>Are you sure that all the drives on this 7190 Model 100
are not in use? Type Y or N:y
>>Test Start ..Done.
```

3. After the diagnostics have completed, the 7190 Model 100 does a web walk. Allow the web walk to complete and check the LED status (refer to *Installation and User's Guide* for the 7190).
4. The diagnostics run is complete.

5.1.6.3 Set/Reset Service Mode

This function places an SSA drive in service mode in anticipation for physically removing the drive from the SSA network. Multiple adjacent drives can be placed into service mode.

Note

Only drives configured in a loop topology or at the end of a string can be set to service mode without disrupting activity on downstream SSA drives. Only the master 7190 Model 100 can set reset service mode.

Service mode is automatically reset if a web reset occurs.

A reset on one drive will automatically reset any other service mode drives.

Use the following procedure to set service mode:

1. Stop all I/O activity to the target drive.
2. Select **Set/Reset Service Mode** from the SSA Service Functions menu. This prompts you for a target drive:

```
>>Input selection ('TxDy, for drive ID x and LUN y):t0d2
```

3. Type in the required information. In our example we have chosen to set t0d2 (This is SCSI ID 2 in Figure 111 on page 142):

```
>>Input selection ('TxDy', for drive ID x and LUN y):t0d2
>>Are you sure that the device is not in use? (Type Y or N):
```

```
>>Cannot locate alternate 7190 model 100 model 100partition, service
>>mode not allowed.
```

   To resolve this condition, create an alternate partition on another drive.

4. If you type **y**, and service mode sets successfully, the message **>>done** appears. The check LED on the target drive is turned on.

   If the operation has failed for one of the following reasons:

   - SSA web is not a loop topology
   - Another disk drive still in Service mode.
   - Drive not at end of string.

   The following message appears:

```
>>The attempt to put the disk drive into Service mode has failed.
```

5. To reset service mode, Select **Set/Reset** mode again and target the same drive.

   This will turn the check LED off.

5.1.6.4 Error Log Analysis

The SSA service functions create an error log file (nnn7190ERR.log) to track all SSA errors. The error log is written to the path specified in the 7190 Model 100 configuration file. The directory (or partition) is created during the installation of the 7190.

The error log analysis option displays the error log and Service Request Numbers for the errors that require attention.

Errors that are marked as old by Clear check mode (hdref refid=clear.), are not reported in the next error log analysis display.

Select **7. Error Log Analysis** from the SSA Service Functions menu.Information is displayed in the following format:

```
>>SSA log analysis:
Mon Nov 30 20:30:03 1997:nnn.T0.D1.AC7E5B59 SRN=2A003
Mon Nov 30 20:30:06 1997:nnn.22000000 SRN=45009

>>
```

The characters **nnn** represents the name of the 7190 Model 100 as specified in the configuration file.

An error that is reported by Async_Alert such as Loss of Redundant Power/Cooling is only logged in the system that contains the SSA Master 7190.

Refer to *Installation and User's guide* for service request number translation.

5.1.6.5 Clear Check Mode

Use this function to clear the check status of a particular drive.

When the SSA service functions log an error for a device, the SSA topology status changes to **CK**. The device might, or might not, still be operational.

In our example, the second drive has a status of **CK**:

```
nnn7190.22000000

SSA.PORT.A1=
=N.T0D0.AC7C1066.OK=
=N.T0D1.AC7E5B59.CK=
=N.T0D2.AC7C0AC5.OK=
=N.T0D3.AC7C0A83.OK=
```

1. Select **8** (″Clear Check Mode) from the SSA Service Functions menu. This option prompts you for the target drive:

```
Clear_check_mode
>>Input selection (TxDy for drive ID and LUN y):t0d2
```

2. Once you select the target drive, the drive will be cleared and the process is complete.

5.1.6.6 Power Checks

DANGER

The product enclosure should be opened only by authorized service personnel. Internal to the product enclosure are exposed areas of high voltage. Incorrect handling of internal components can cause harmful electric shock.

1. Verify that the power switch is off.

   DANGER

   An electrical outlet that is not correctly wired could place hazardous voltage on metal parts of the system or the products that attach the system. Verify that the wiring and grounding of the electrical outlets you use have been checked by a licensed electrician.

2. Verify that AC power is present at the electrical outlet being used.
3. Verify that the 7190 Model 100 power cord is fully inserted o n the mainline side and on the power connector side of the 7190.
4. Set the 7190 Model 100 power switch to off.
5. Remove the fuse and check whether it is blown. If the fuse is blown, replace it with a 250 VAC fast blow fuse. For a part number, refer *Installation and User's Guide.*
6. Set the 7190 Model 100 power switch on. Verify that the power LED is on.
7. If the power problem is corrected, return the procedure that sent you here.

If your power problem is not corrected, contact your service representative and report the problem.

5.1.6.7 Hosts System Checks

A host problem may be caused by, but not limited by the following conditions:

- A device may have been plugged in while the host was powered on.
- There might be an SSA cable failure.
- The host system may have been shut down incorrectly.

Use the following procedure to troubleshoot host system problems:

1. Check the SCSI cable that connects the 7190 Model 100 to the host SCSI adapter for pin damage. Make sure that the cable is of high quality, taking into account the SCSI cable specifications.
2. Verify that the SSA cables attached to port 1 and port 2 are correctly connected. Check the quality of the cables.

3. Verify that SSA devices have not interrupted the operation of the SSA bus. SSA devices can lockup whenever the normal operation of the SSA bus has been interrupted and the host system does not properly remove outstanding requests:
   - Set the power switch to off on the host system.
   - Power cycle both the 7190 Model 100 and the SSA devices attached
   - Set the host system power switch to on.
   - Restart the host system operating system.

   Cycling power to the drives and the 7190 model 100 model 100clear lockup conditions that the outstanding requests may have caused.
4. Test the host SCSI port by attaching a known working device, such as a SCSI cd rom, to the port and checking whether it operates correctly. If you are unable to communicate with the device, the problem is most likely not the 7190 Model 100, but a system hardware or configuration problem.

5.1.6.8 Status LED Checks

Refer to Figure 10 on page 17 for this procedure.

5.1.6.9 Collecting system information

When you call a service representative, you will be required to provide certain information before your problem can be processed. We advise you to collect this information before you start your own problem-solving procedure. The information you record will be used in the problem determination procedure that we provide.

The following is a basic template of the information that you should document:

- Customer name
- Contact person and telephone number
- Host type and model number
- Operating system type and revision level
- SCSI adapter model number/s
- 7190 Model 100 serial number
- 7190 Model 100 vital product data (VPD)
- SSA Storage Subsystem model number and serial number
- Description and configuration of SSA topology information such as SSA loop structure and number of drives (“Show SSA Topology” on page 127)
- Record of file subsystem information such as mirroring and other raid levels
- Description of problem
- LED POST pattern (5.1.6.8, “Status LED Checks”)
- If an individual drive was found defective through a message on the host monitor or a recorded SCSI error status in the Operating System log, record the failed SCSI/LUN as seen from the host

5.1.7 Problem Determination Procedure

Your system can be categorized into two components, the host computer system, and the storage subsystem. These two components are typically supported by separate field service organizations. The host computer system consists of the hardware, the operating system, and device drivers. The storage subsystem consists of a 7190 model 100 Model connecting to a 713x IBM SSA disk subsystem.

The objective of the problem determination procedure is to:

- Categorize any problems as either host computer problems, or storage subsystem problems.
- Identify and fix all problems related to the storage subsystem.
- Identify and direct all the host computer related problems to the appropriate field support personnel

The problem determination procedure may require one or more disks to be power cycled. Make sure that I/O is halted before undertaking the procedure and that you have a full backup of the storage attached to the SSA network.

We have mapped the problem determination procedure into the flowcharts in Figure 112 and Figure 113 on page 150 Figure 112 represents the main routine of the problem determination procedure.

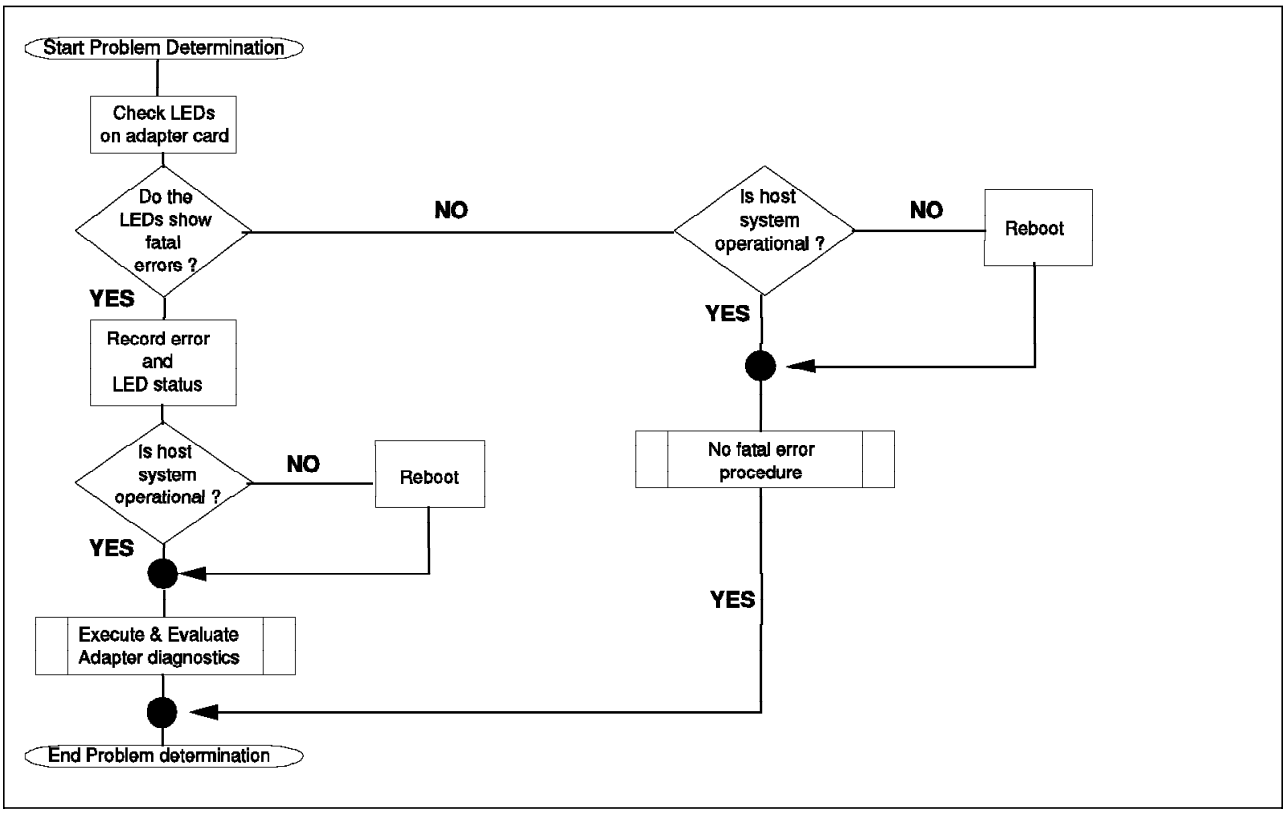

Figure 112. The Main Troubleshooting Procedure

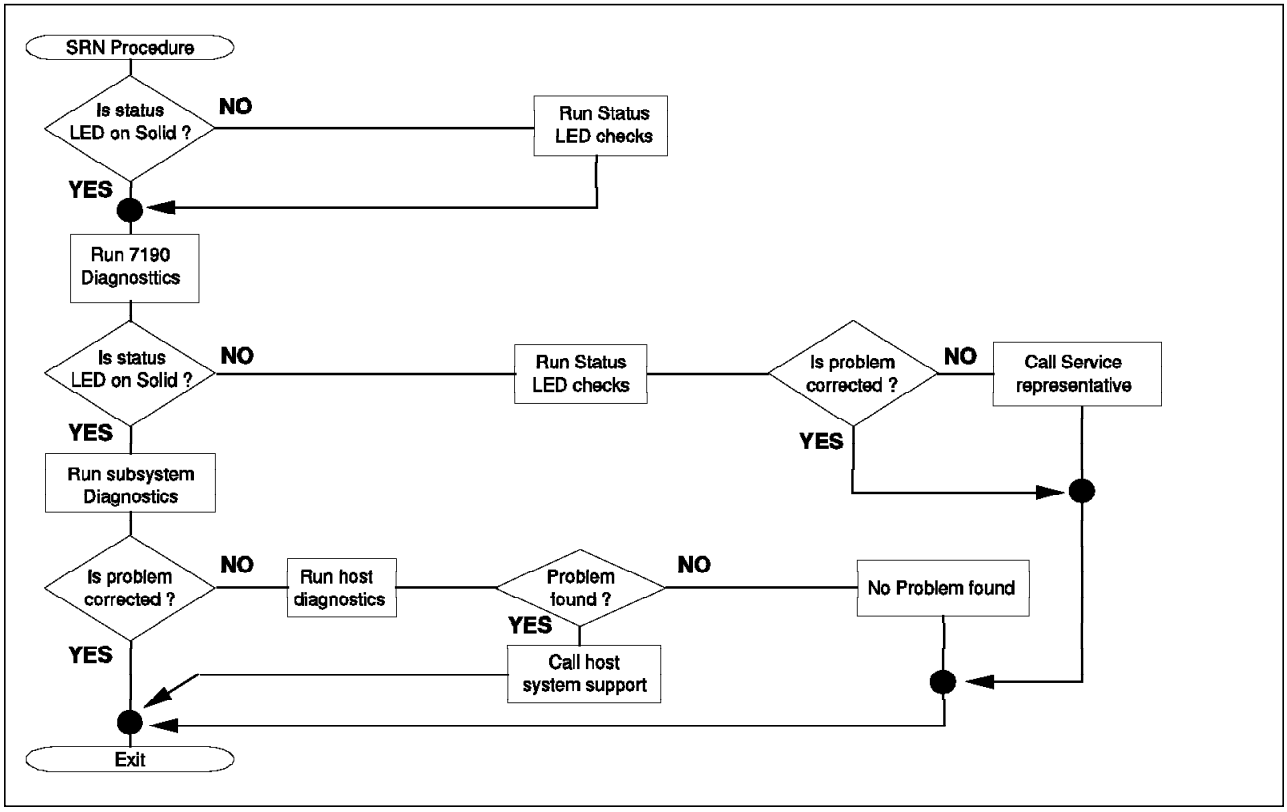

Figure 113. The Procedure to Follow if no SRNs are Reported in the Log

The problem determination procedure represented by the flowcharts above, is described as follows:

1. The first step is to confirm that you have the SSA Service Function up and running. The steps that follow rely on functions available through the SSA Service Function. If the service functions are running, you can type ? at any time to view your options. If the function is running, proceed to step 5.
2. Proceed with this step if you answered no to the question posed in Step 1. Confirm that the package is installed by typing **./SSA_SERV - dnnn** (Replace **nnn** with the name of the 7190). If the package is already installed, proceed to Step 4.
3. Proceed with this step if you answered no to both Step 1 and Step 2. Install the SSA Service functions in order to copy the functions that are required to continue the problem determination procedure to. Follow the installation procedure in " Installing Services Function on Sun" on page 117.
4. Proceed with this step if you have completed Step 3. Start the SSA Service Functions as described in 5.1, "7190 Model 100 operations" on page 125 and continue with step 5.
5. Run an error log analysis as described in 5.1.6.4, "Error Log Analysis" on page 146 to check whether any SRN numbers have been generated by the SSAU health check. You will be asked for the SRN numbers if you call an IBM service representative. You will also use the SRNs further in the problem determination procedure. If an SRNs is reported, you have a storage subsystem problem and should refer to the service guide for the SSA storage subsystem you are using.

6. Run this step if you don't have an SRN number after completing step 5 (Figure 111 on page 142). Check whether the power LED on the 7190 Model 100 is on to confirm that it has power. If the power LED is on, skip to Step 7. If the power LED is off, run the power-checking procedure described in 5.2.6.6, "The Troubleshooting Procedure" on page 179 and continue with Step 7 if the power problem is corrected. If you are unable to correct a power problem, call your service representative.
7. Check whether the status LED is on steadily or whether it is flashing. If it is flashing, go to status LED checks as described in 5.2.6.3, " Adapter LED Analysis" on page 178 and correct the problem. If you are unable to stop the LED from flashing, call your service representative.
8. Proceed with this step if you answered yes to the question posed in Step 7. Run the 7190 Model 100 diagnostics as described in 5.1.6.2, "7190 Model 100 Diagnostics" on page 144. This procedure will cause the 7190 Model 100 to do a reset, which may resolve the problem. If the problem is not corrected, record the system information (5.1.6.9, "Collecting system information" on page 148) and call your service representative.
9. Check whether the status LED is on steadily (not flashing). If not, run the LED checks and continue with the next step if the problem is resolved. If the status LED is not on steadily and you are unable to correct the problem, record the system information (5.1.6.9, "Collecting system information" on page 148) and call your service representative.
10. Run the subsystem diagnostic test by setting the 7190 Model 100 to mode 15. The subsystem diagnostics will run a test on all the drives attached to the 7190. The diagnostics will return a status code if an error is detected on one of the drives, in which case you will have to take the appropriate action, whether it be a swapout of the failed drive or a call to your service representative. If no drive failures are detected, or you correct the problem, proceed to the next step.
11. Run the host system diagnostics described in 5.1.6.7, "Hosts System Checks" on page 147 to verify whether or not the problem is host system related. If a problem is found, record the system information (5.1.6.9, "Collecting system information" on page 148) and call your service representative.
12. At this stage, you should have been able to isolate and correct the problem. If a problem still exists, record the system information (5.1.6.9, "Collecting system information" on page 148) and call your service representative.

5.2 IBM SSA Interface Controller for Sun SBus Operations

The SSA Service Aid Utility is used to monitor the the SSA network and its components, maintain the SSA web, add and replace devices, upgrade the SSA Interface Controller and SSA disk drives, and troubleshoot the SSA network.

The utility is run from a UNIX command line and requires root or equivalent privileges. The command syntax is as follows:

```
ssau -(option)
```

You can access online help on the **ssau** by typing:

```
# ssau -h
```

Online help is divided into adapter operational help, health check operation help, disk operation help, and error log analysis help: :

- Adapter operational help **ssau -h a**
- Health check operational help **ssau -h c**
- Disk operational help **ssau -h d**
- Error log analysis help **ssau -h e**

5.2.1 Monitoring the SSA Network

This section explains how to collect SSA network information that is vital to maintain the network. We recommend that you file the output of the commands in this section.

5.2.1.1 SSA Adapter Operations

To display the options type:

ssau -h a

Adapter Specify (0|1) You must direct your command to a controller 0 or 1.

- Flashes green LED of specified adapter
- Default is 5 s

command line syntax:

ssau -a 0 -i (identify adapter 0 for the 5 s)

ssau -a 1 -i 10 (identify adapter for 10 s)

Display Vital Product Data (VPD) (-v)

- Displays hardware or firmware revision
- Displays adapter memory size

command line syntax:

ssau -a 0 -v (display adapter 0 vital product data)

Display the SSA web map (-w)

- Shows the logical connection of the nodes to the adapter
- Displays the node UIDs
- Shows the node positions

command line syntax:

ssau -a 0 -v (display adapter 0 product data)

5.2.1.2 Disk Operations

To display the options type:

ssau -h d

All SSAU commands have to be directed to a specific controller.

Disk drive identity:

- Flashes LED of specified disk
- Default duration is 20 s

  command line syntax:

  ssau -d •a (n) -1 (duration) (identify disk drive n on adapter a)

  For example,

  ssau -6 6 -i (identify disk drive 6 of adapter 0 for 20 s)

  ssau -d 1:7 -i 30 (identify disk drive 7 of adapter 1 for 30 s)

Display disk drive Vital Product Data (VPD) (-v): This is use to verify the following information:

- Vendor Id
- Product type
- Model
- ROM revision
- RAM revision
- Serial number
- Microcode load
- Plant of manufacture
- Date of manufacture
- FRU count
- FRU length
- Assembly part number
- Assembly EC level
- Card assembly part number
- Card assembly EC level

command line syntax:

ssau -d •a (n) -v (display VPD of disk drive n on adapter a)

For example,

ssau -d 3 -v (display VPD of disk drive 3 on adapter 0)

ssau -d 1:8 -v (display VPD of disk 8 on adapter 1)

5.2.2 Maintaining the SSA Web

You use the SSCF utility, which is copied onto the system diskette during the installation, to maintain the SSA web. It is essential that you back up your data and make a copy of the configuration map before you use the utility as incorrect usage of the utility can result in data loss:

command line syntax:

sscf {-a 0|1} {-c} {-d} {-e} {-h} {-m mapfile} {-p} {-r} {-u} {-y

Two or more options can be combined in any order. The priority of the operation is:

-m or -c, -u, -r, -d, -pe or -p

- **a 0|1** Specify the target adapter for the command i
- **c** Create a new map of alL SSA devices. This option forces assignments for all disk devices on the web that the driver detects.

  Attention

  This option may result in reassignment of SCSI IDs to SSA UIDs. The option should only be used by administrators who have a thorough knowledge of Solaris and how the user applications administrators.

- **-d** Display the current SSA UID to SCSI ID configuration map
- **-e** Erase persistent configuration map; use only with the **-p** option
- **-m** Map: specify a **map file** which will be used to set the map
- **-p** Persistent: Save the current web configuration as the default at boot
- **-r** Replace drive: notify the driver that the SCSI drive removed will be reassigned to the next drive added.
- **-u Update** current map by assigning SCSI IDs to all known devices that do not yet have SCSI ID assignments.

  Note

  This option may result in reassignment of SCSI IDs to SSA UIDs and it should only be used by administrators.

- **-y** Yes mode: do not display confirmation prompts for the **-c, -m, -p** and **-u** options. Instead, the command proceeds as if the operator had answered yes to each question. This option is useful for inclusion in shell scripts.

Figure 114 on page 155 is a logical representation of an SSA adapter.

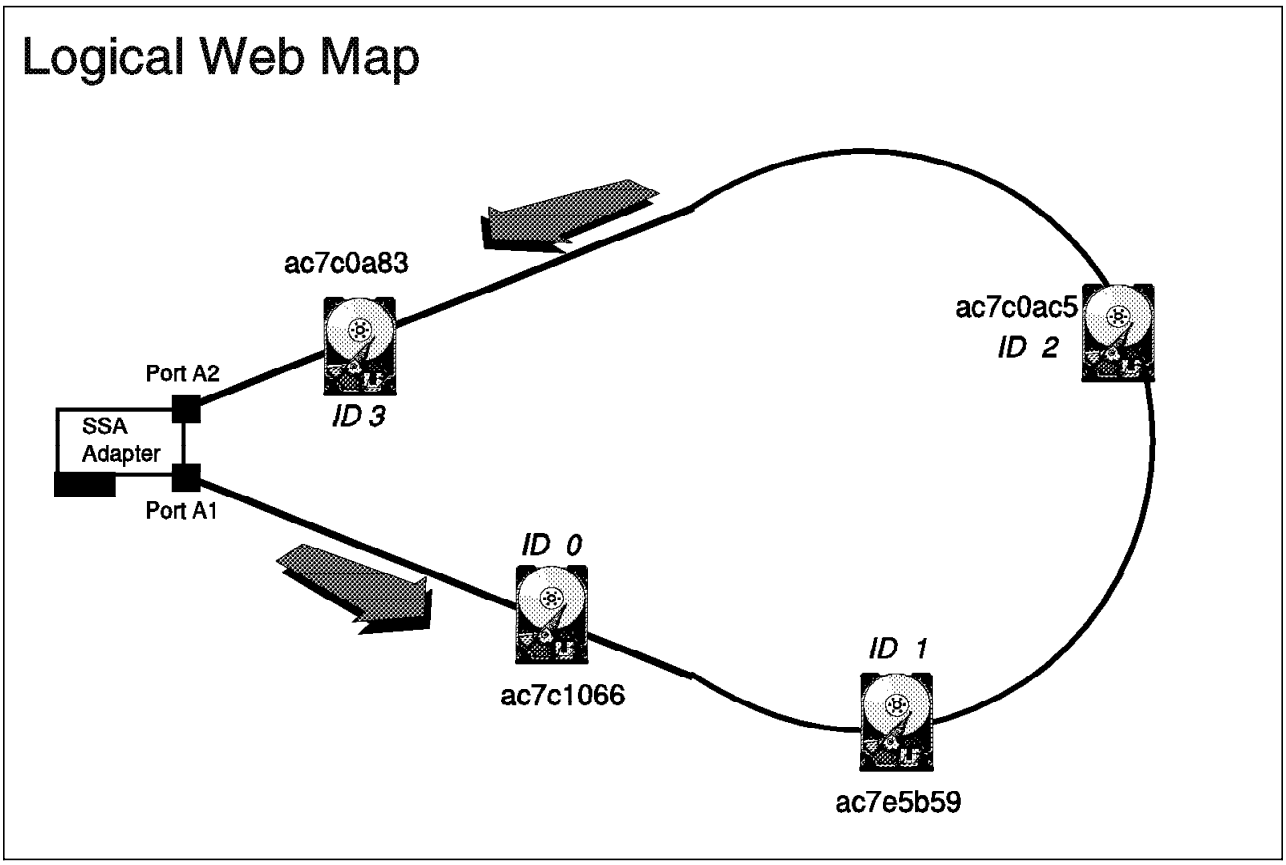

Figure 114. Representation of a Logical Web Using an SSA Adapter

5.2.2.1 Display the Configuration Map

- Displays adapters configuration map
- Displays UIDs and assigned SCSI IDs

command line example:

sscf -a 0 -d

The output for the above command after a fresh installation is:

```
root> sscf -a 0 -d
# IDs of installed adapters: a b
# Configuration map for adapter 0
# UID          UID        SCSI
# High   Low        ID
  0004   ac7c1066   00
  0004   ac7e5b59   01
  0004   ac7c0ac5   02
  0004   ac7c0a83   03
```

As in Figure 114, note the SCSI ID assignment order in an IBM SSA Interface Controller for Sun SBus loop topology. The SCSI IDs are assigned in increments according to the device position (in node hop counts) relative to port 1.

5.2.2.2 Updating the Configuration Map (-u)

This option is use to update the current configuration map by assigning SCSI IDs to all unassigned devices that are recognized by the controller. The lowest available SCSI ID is assigned to the unassigned device closest to port 1. Successively higher available SCSI IDs are assigned to subsequent unassigned devices.

command line example:

```
# sscf -a 0 -u
Ids of installed Adapters: 0 1
Continuing will update the map for Adapter0: Proceed (y/n) y
The configuration map has been modified.
#
```

```
root> sscf -a 0 -d
# IDs of installed adapters: 0 1
# Configuration map for adapter 0
# UID         UID        SCSI
# High   Low        ID
  0004   ac7c1066   00
  0004   ac7e5b59   01
  0004   ac7c0ac5   02
  0004   ac7c0a83   03
  0004   ac7e5a9b   04
  0004   ac7c0ac5   05
```

Existing device IDs are reserved for their assigned devices. The next available SCSI ID (SCSI ID 3 in this example) is assigned for the newly recognized device that is closer to port 1. The subsequent devices that are newly recognized are assigned successive available Ids.

The **-p** option **must** be used after the update to make the map persistent:

```
:
root>
root> sscf -a 0 -p
Ids of installed Adapters: 0 1
About to modify file '/kernel/drv/sunlight.conf'! Proceed? (y/n) y
Backup file '/kernel/drv/sunlight.conf.bak' has been written.
The drivers configuration file '/kernel/drv/sunlight.conf' has been
modified.
root>
```

Alternatively, the **-u** and **-p** options can be combined like this:

```
root>
root> sscf -a 0 -p
IDs of installed Adapters: 0 1
Continuing will update the map for Adapter 0! Proceed? (y/n) y
About to modify file '/kernel/drv/sunlight.conf'! Proceed? (y/n) y
Backup file '/kernel/drv/sunlight.conf.bak' has been written.
The drivers configuration file '/kernel/drv/sunlight.conf' has been
modified.
root>
```

5.2.2.3 Creating a New Configuration Map

c uses the default algorithm for assigning SCSI IDs as discussed in hdref refid=sbdcmap. and can result in the reassignment of SCSI IDs. If any programs rely on specific devices attached to the IBM SSA Interface Controller for Sun SBus, they may need to be reconfigured according to the newly assigned SCSI IDs.

The option can be used in the following scenarios:

- The adapter's web consists of all new devices that have been installed while the system is up and running.
- New drives have been added and existing drives are not been accessed.
- SCSI ID to UID mapping information is documented and you intend to redirect programs to the same physical devices with new SCSI IDs. (This option is advisable only for Solaris experts.)

Create a new configuration map as in the following command line example:

```
root>
root> sscf -a 0 -c
IDs of installed Adapters: 0 1
Continuing will create a default map for Adapter0: Proceed (y/n)
y
The configuration map has been modified
root>
```

Verify the new map as follows:

```
root> sscf -a 0 -d
# IDs of installed adapters: 0 1
# Configuration map for adapter 0
# UID         UID        SCSI
# High   Low        ID
  0004   ac7e5a96   00
  0004   ac7c1066   01
  0004   ac7e5b59   02
  0004   ac7c0a83   03
  0004   ac7c0ac5   04
  0004   ac7c0bf1   05
```

You can see by comparing the display output to the diagram that the SCSI IDs have been assigned according to node position relative to port 1.

Use the **-p** option to make the new configuration persistent:

```
:
root>
root> sscf -a 0 -p
IDs of installed Adapters: 0 1
About to modify file '/kernel/drv/sunlight.conf'!  Proceed? (y/n) y
Backup file '/kernel/drv/sunlight.conf.bak' has been written.
The drivers configuration file '/kernel/drv/sunlight.conf' has been
modified.
root>
```

Alternatively, the **-c** and**-p** options can be combined as follows:

```
:
root>
root> sscf -a 0 -c -p
IDs of installed Adapters: 0 1
Continuing will update the map for Adapter 0! Proceed? (y/n) y
The configuration map has been modified.
About to modify file '/kernel/drv/sunlight.conf'!  Proceed? (y/n) y
Backup file '/kernel/drv/sunlight.conf.bak' has been written.
The drivers configuration file '/kernel/drv/sunlight.conf' has been
modified.
root>
```

Note that the first access to a new device being powered up can cause the Solaris target driver to display error messages because the device is not yet up to speed.

5.2.2.4 Creating and Using a Map File

A backup ASCII file can be created by redirecting the output of the **sscf -d** option. The file contains associations between SCSI IDs and SSA UIDs and is used to override existing configuration maps, thus serving as a useful recovery tool for a corrupted or missing map.

Command line example:

```
root>
root> sscf -a 0 -d > mapfile
root>
```

The contents of the map ile in our example is:

```
root> sscf -a 0 -d
# IDs of installed adapters: 0 1
# Configuration map for adapter 0
# UID          UID        SCSI
# High   Low        ID
  0004   ac7e5a96   00
  0004   ac7c1066   01
  0004   ac7e5b59   02
  0004   ac7c0a83   03
  0004   ac7c0ac5   04
  0004   ac7c0bf1   05
```

You can modify the map file to change the SCSI ID to UID associations by using a text editor like **vi** or **emacs** Any modification you make to the map file have to be consistent with the physical positions of SSA devices as incorrect usage can result in data loss.

Use the **-m** option to use the initiate the associations in the map file:

```
root>
root> sscf -a 0 -m mapfile
IDs of installed adapters: 0 1
Continuing will apply the map for Adapter 0! Proceed? (y/n) y
The configuration map has been modified.
eoot>
```

Use the **-p** option to make the new configuration persistent:

```
root>
root> sscf -a 0 -p
Ids of installed Adapters: 0 1
About to modify file '/kernel/drv/sunlight.conf'! Proceed? (y/n) y
Backup file '/kernel/drv/sunlight.conf.bak' has been written.
The drivers configuration file '/kernel/drv/sunlight.conf' has been
modified.
root>
```

Alternatively, use the **-m** and **-p** options combined:

```
:
root>
root> sscf -a 0 -p mapfile
IDs of installed Adapters: 0 1
Continuing will update the map for Adapter 0! Proceed? (y/n) y
The configuration map has been modified.
About to modify file '/kernel/drv/sunlight.conf'! Proceed? (y/n) y
Backup file '/kernel/drv/sunlight.conf.bak' has been written.
The drivers configuration file '/kernel/drv/sunlight.conf' has been
modified.
root>
```

Note: that the file name used for the map file is your choice.

5.2.3 Adding Drives Devices

This section deals with the expansion of IBM SSA Interface Controller for Sun SBus networks, device removal, and device replacement. It is important that you follow the exact procedures to avoid data loss. Make sure you have a copy of the map file.

5.2.3.1 Adding Drives to an SSA Interface Controller SSA loop

This section requires you to make a single break in the loop. Multiple breaks in the loop can isolate drives from the IBM SSA Interface Controller for Sun SBus, thus resulting in loss of data access. If multiple drives are to be added, we suggest that you add them to the same break in the loop. If you require the drives to added at separate breaks, you will have to follow this procedure for each break.

Objective: add a new drive to the loop between two existing drives

1. Identify the place in the SSA loop where the drives are to be inserted (we have chosen to insert the drive between ***C*** and ***D***). Break the SSA loop at this spot by disconnecting the cables linking the two drives. This will divide the loop into two strings (Figure 115) The steps are as follows:.

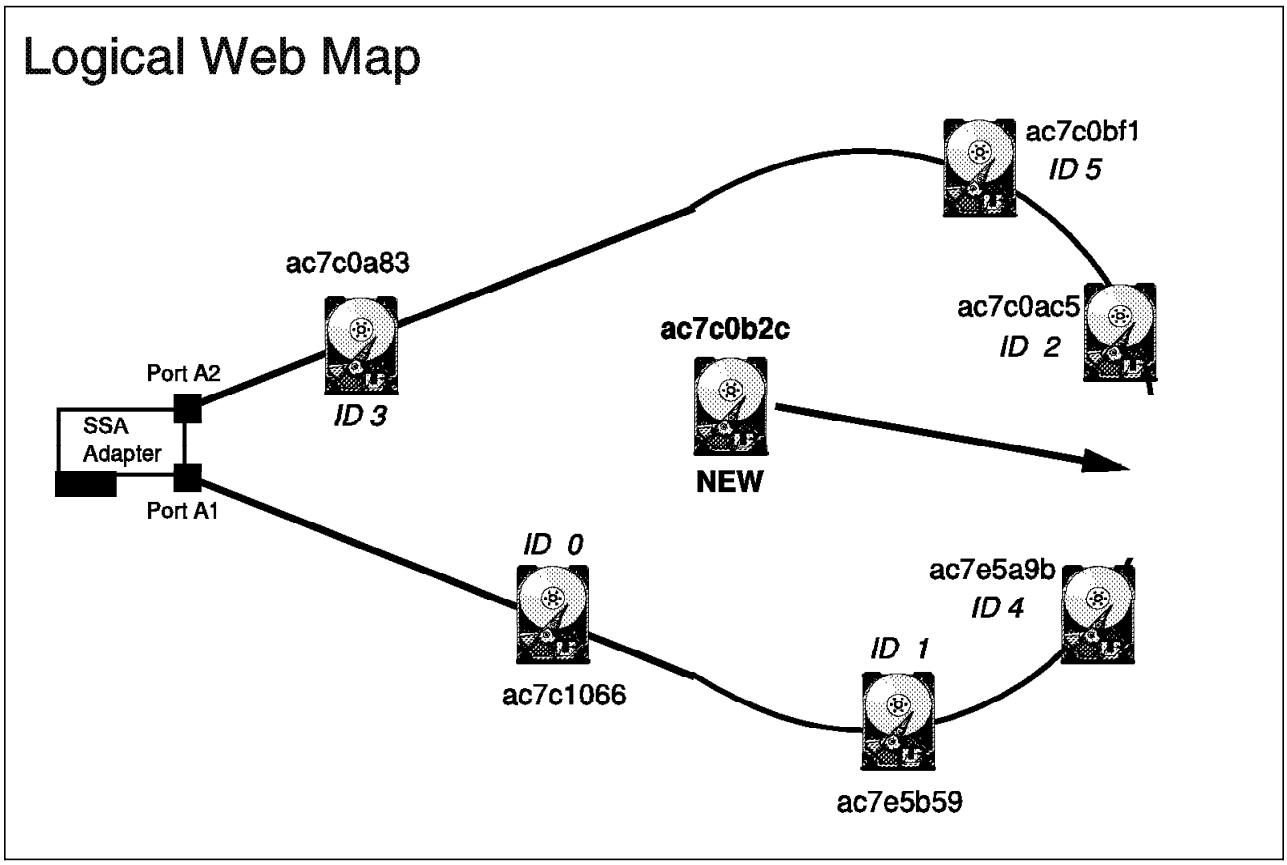

Figure 115. Inserting a New Drive into an SSA Interface Controller Loop

a. Attach the drive to one of the strings and close the loop.

b. Power up the drives.

c. Check the drive to make sure the green LEDs are on and the yellow LED is off.

d. Assign SCSI IDs to the new drives with the following command:

   sscf -a 0 -u -p (for drives connected to adapter 0)

   or

   sscf -a 1 -u -p (for drives connected to adapter 1)

e. If additional drives are to be added, repeat Steps 1-5.

f. Reboot Solaris:

   reboot -- -r (from Solaris prompt)

   or:

   boot -r (from open boot prompt)

g. Figure 116 represents the procedure result.

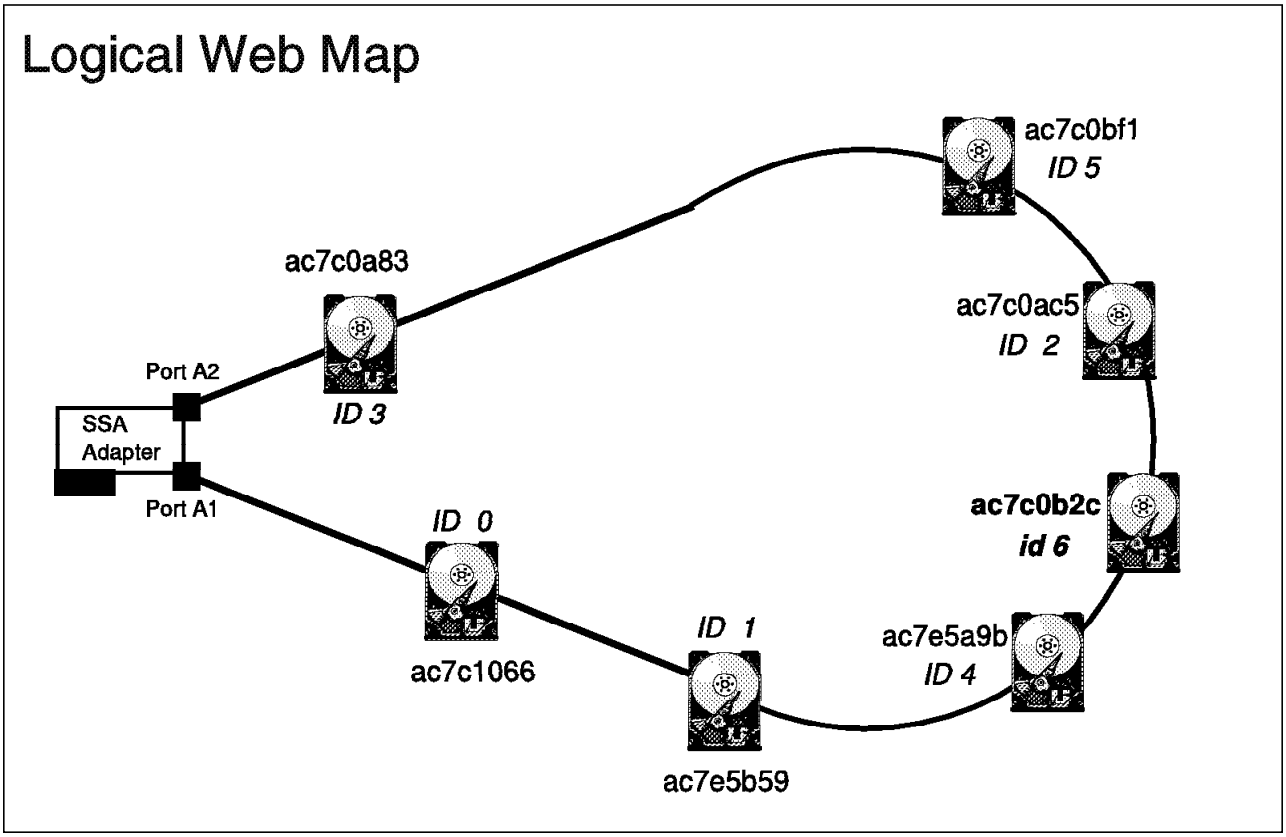

Figure 116. The Newly Added Drive in the SSA Adapter Loop

h. Display the configuration map and compare it with Figure 116:

```
root> sscf -a 0 -d
# IDs of installed adapters: 0 1
# Configuration map for adapter 0
# UID          UID        SCSI
# High   Low        ID
  0004   ac7e5a96   00
  0004   ac7c1066   01
  0004   ac7e5b59   02
  0004   ac7c0a83   03
  0004   ac7c0ac5   04
  0004   ac7c0bf1   05
  0004   ac7c0b2c   06
```

i. The new drive is ready for use.

5.2.3.2 Adding Drives to a String

We recommend that you add new drives to the end of a string as drives cannot be inserted between nodes on an active string without interrupting access to drives that are downstream from the point of insertion. If you must insert drives between existing nodes, there is a separate procedure. Make sure that you have a copy of the web file.

Adding new drives to the end of a string:

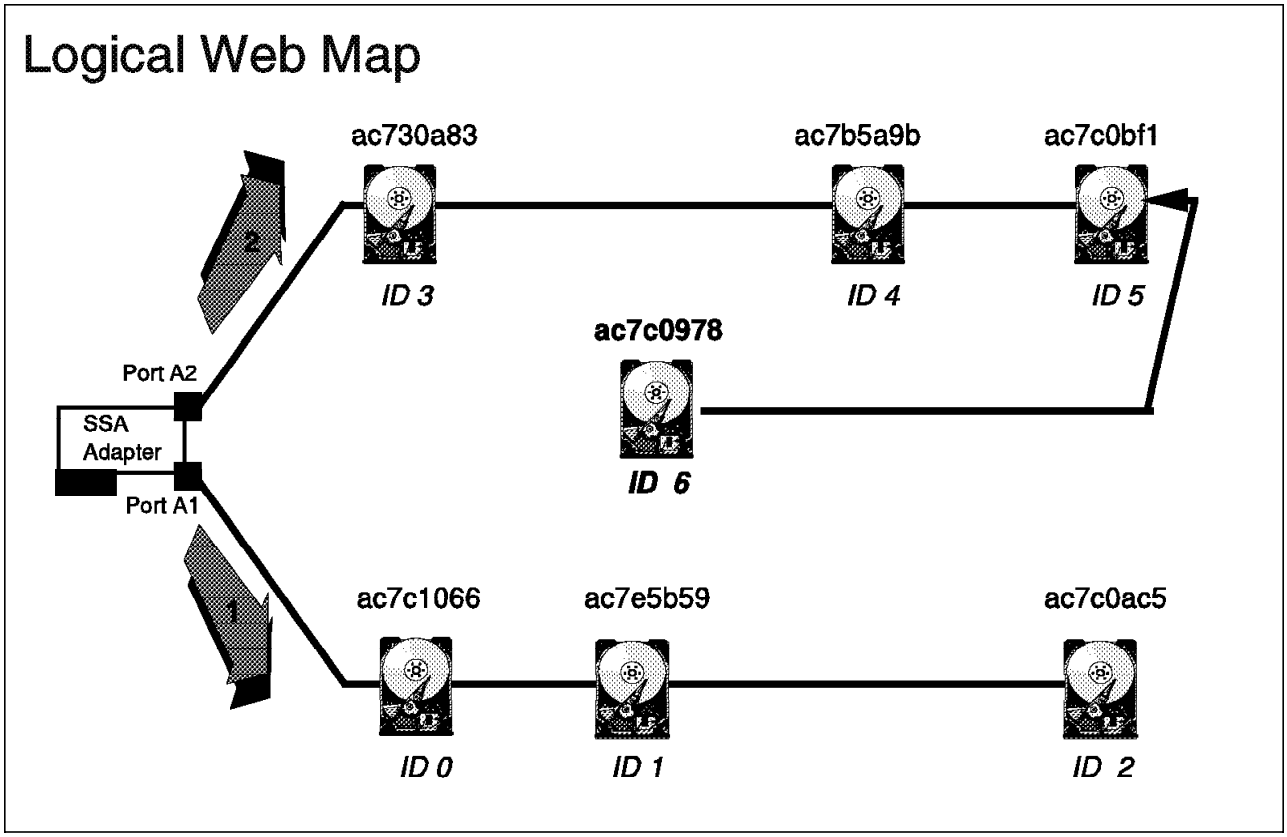

Figure 117. IBM SSA Interface Controller in a string topology.

1. Identify which string you will add the drive to. We have chosen string 2, which originates from port 2. We will extend the cable from the drive located at the end of the string (SCSI ID 5).

2. Cable the new drives to the end of the string. We have added the drive with UID ac7c0978 (see the spare drive in Figure 117)
3. Power up the new drives.
4. Check each drive, except the last, to make sure that there two green LEDs are on constantly and their yellow LEDs are off. The last device should have one green light on constantly and the other one blinking.
5. Assign SCSI IDs to the new drives:

   sscf -a 0 -u -p (for drives connected to controller port 0)

   or:

   sscf -a 1 -u -p (for drives connected to controller port 1)
6. Shut down Solaris.
7. Reboot Solaris specifying the reconfiguration option:

   reboot -- -r (to reboot from the Solaris prompt)

   or:

   boot -r (to reboot from OpenBoot prompt)

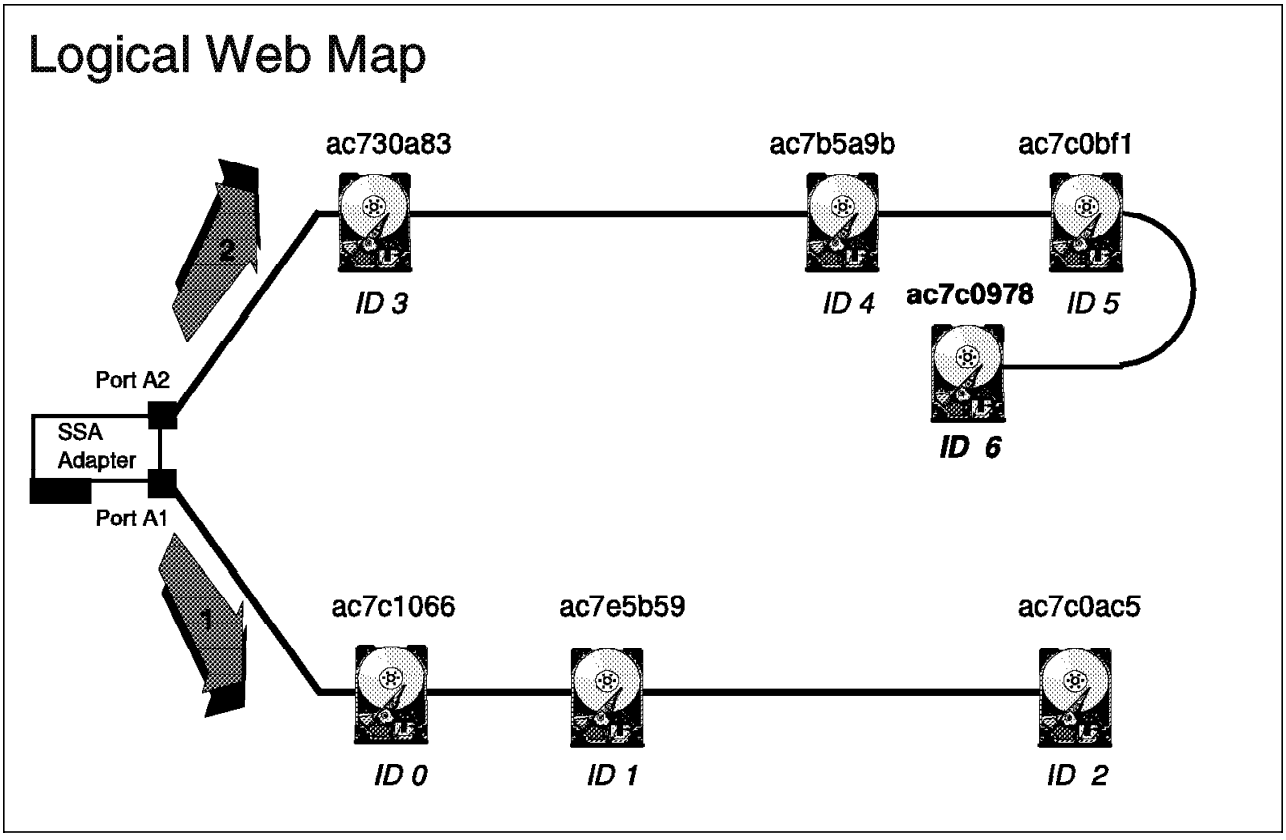

Figure 118. The Completed Addition Procedure

Display the configuration map to confirm the procedure and compare with Figure 118:

```
root>
root> sscf -a 0 -d
# IDs of installed adapters
# Configuration map for controller 0
# UID         UID             SCSI
# High   Low             ID
  0004   ac7c1066        00
  0004   ac7e5b59        01
  0004   ac7c0ac5        02
  0004   ac730a83        03
  0004   ac7b5a9b        04
  0004   ac7c0bf1        05
  0004   ac7c0978        06
root>
```

8. The new drives are ready for use.

Inserting drives between existing nodes:

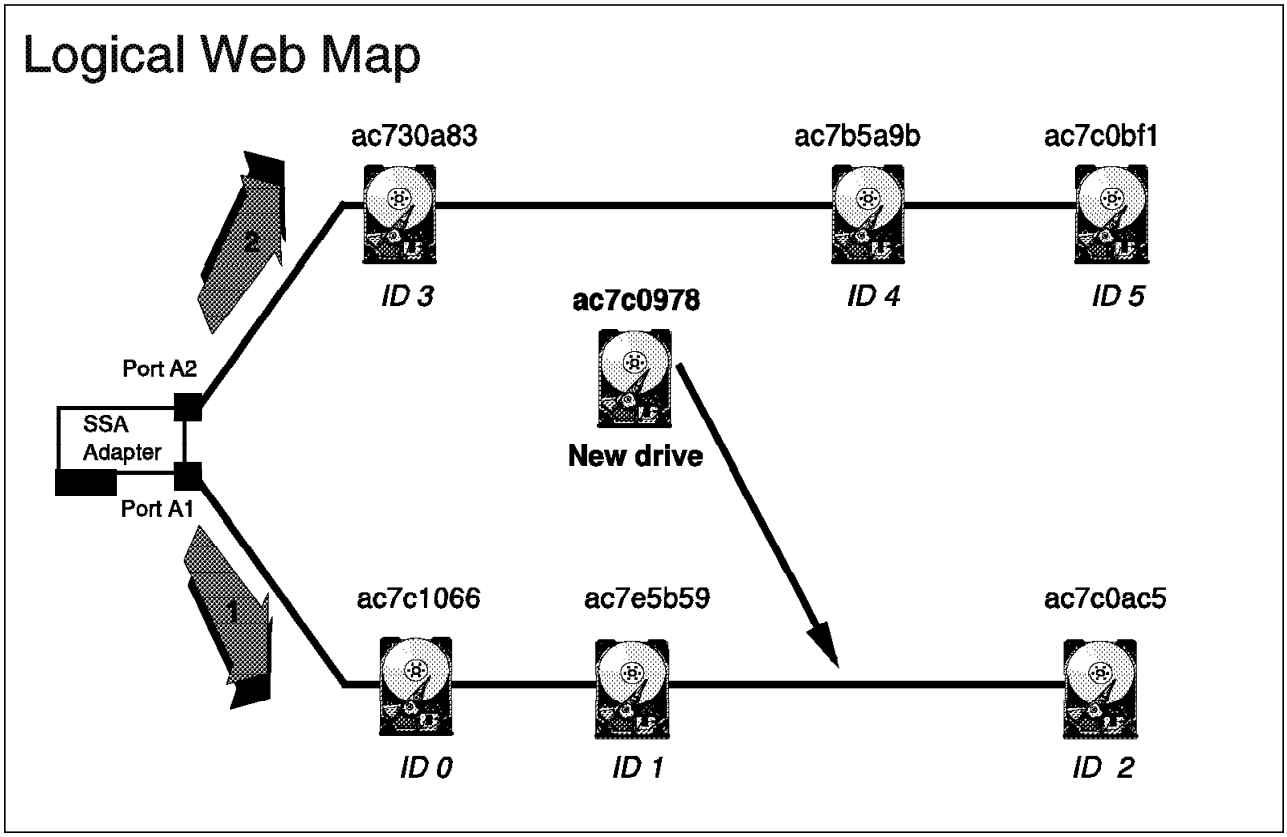

Figure 119. New Drive Being Attached Between Nodes

1. Identify which string and between what two nodes you wish to insert the drive. For our example, we have chosen to insert the drive on string 1 between ***ID 1*** and ***ID 2***

2. Shut down Solaris and power down the system.

3. Break the string at the chosen position and attach the new drives. Attach the drive making sure that the string is continuous and is not split into sections.

4. Power up the new drive.
5. Check the drive to make sure that both of its green LEDs are on and its yellow LED is off.
6. Power up the server and boot the system.
7. Assign SCSI IDs to the new drives:

   sscf -a 0 -u p (for drives connected to adapter 0)

   or

   sscf -a 1 -u -p (for drives connected to adapter 1)
8. Shut down Solaris.
9. Reboot Solaris in reconfiguration mode:

   reboot -- -r (to reboot from the Solaris prompt)

   or:

   boot -r (to reboot from OpenBoot prompt)

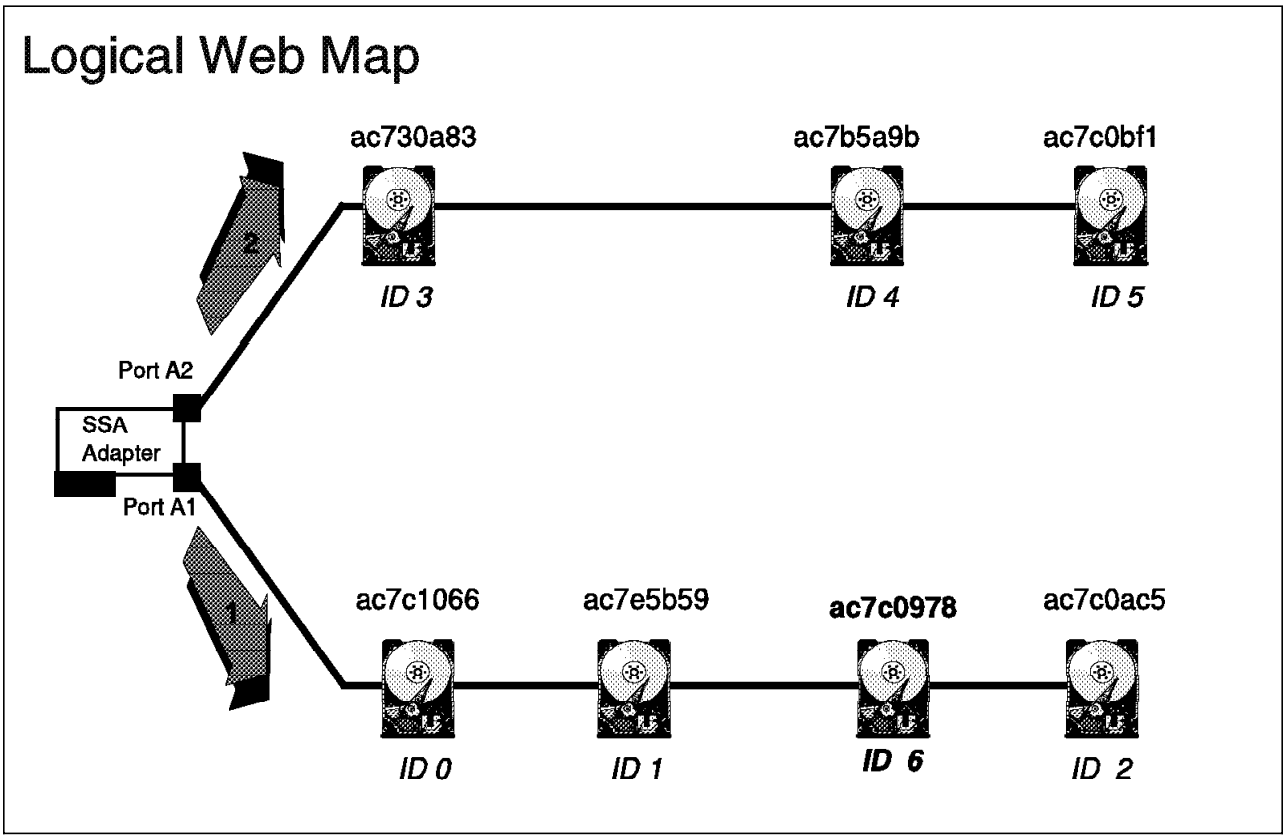

Figure 120. Result Upon Completion of Adding a Drive Between Nodes

Display the configuration map to confirm the result, and compare with Figure 120.

```
root>
root> sscf -a 0 -d
# IDs of installed adapters
# Configuration map for controller 0
# UID         UID             SCSI
# High   Low             ID
  0004   ac7c1066        00
  0004   ac7e5b59        01
  0004   ac7c0ac5        02
  0004   ac730a83        03
  0004   ac7b5a9b        04
  0004   ac7c0bf1        05
  0004   ac7c0978        06
root>
```

10. The new drives are ready for use.

5.2.4 Replacing Devices

5.2.4.1 Replacing the IBM SSA Interface Controller for Sun SBus

THe SSCF utility is not linked to a specific IBM SSA Interface Controller as it affects only the driver. When the system boots up, the driver attaches to the adapter and queries it for its current configuration.

We recommend that you replace the adapter in the same SBus slot as the previous adapter to avoid any reconfiguration. At this level, all that matters is that the card is in the same slot. The adapter UID is not taken into account. If an adapter is added to a new SBus slot, Solaris recognizes it as a new device and creates a new instance for it.The system will have to be rebooted in configuration mode to access the new instance.

The procedure is therefore as follows:

1. Do a clean shutdown of the system and turn the power off.
2. Take antistatic precautions while opening the machine and replacing the faulty adapter with the new adapter.
3. Power up the system and display the web map to verify that the card is operating.

The only difference reported is the UID of the adapter reported in the web map. If the previous adapter was set as the SSA master, you will have to set the master accordingly.

We recommend that you limit the IBM SSA Interface Controller for Sun SBus to the first SBus slots.

There are separate procedures for replacing drives on a loop and replacing drives on a string. Both procedures use the **sscf** and **ssau** utility.

Note

Only one drive can be replaced at a time on an active system.

The first access to a new device being powered up may cause the Solaris target drive to display error messages because the device has not spun up yet. A subsequent access should correct the error by spinning the drive up.

5.2.4.2 Replacing Drives in an IBM SSA Interface Controller Loop

When you replace a drive in a loop, the loop will be broken into two strings, thus allowing continued I/O between the other SSA devices, as shown in Figure 121.

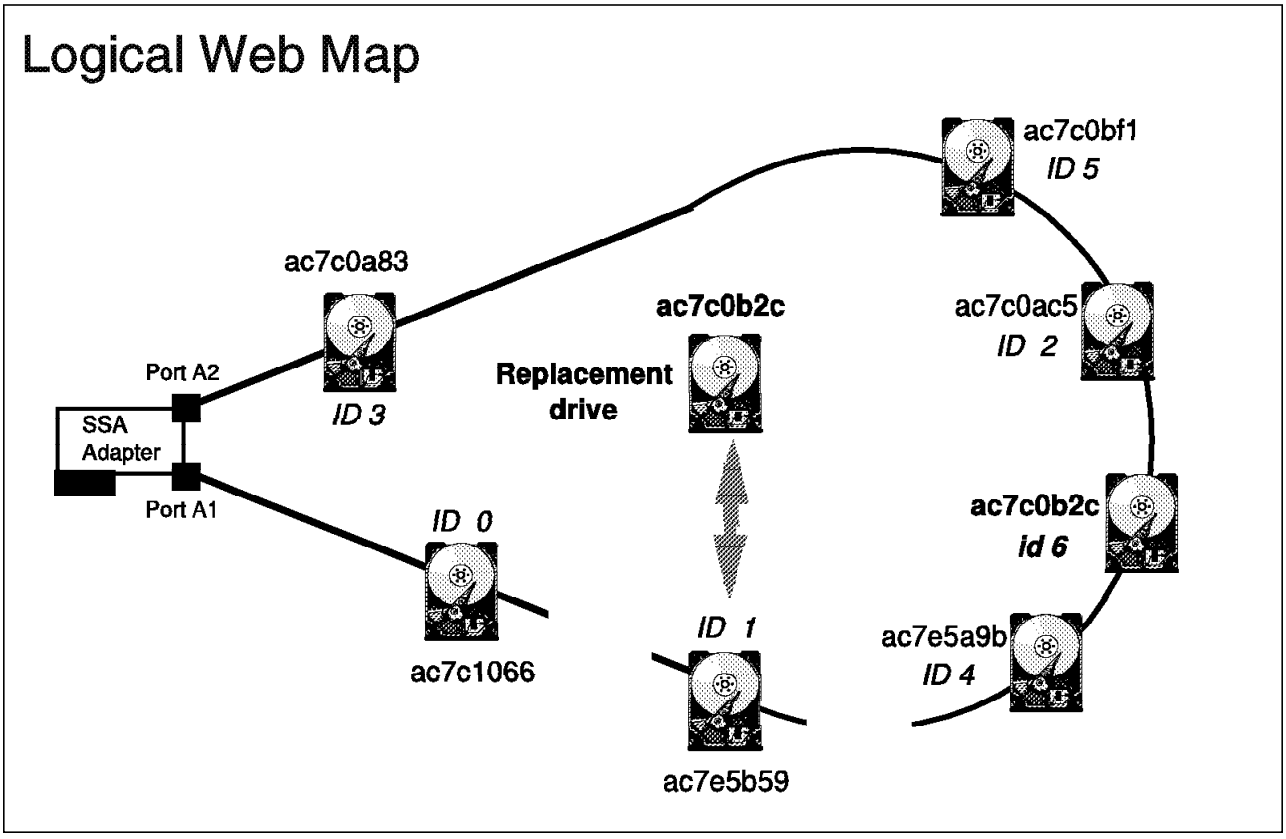

Figure 121. SSA Drive Replacement in a Loop

To replace a drive in a loop, follow this procedure:

1. Identify the drive to be replaced. We have chosen ***ID 1*** from the diagram.
2. Reserve the SCSI ID of the drive being removed for the replacement:

   sscf -a 0 -r (for adapter 0)

   or

   sscf -a 1 -r (for adapter 1)
3. Put the drive that is to be replaced in service mode:

   ssau -d a:n -s set (where **a** is the adapter number and **n** is the SCSI ID of the target drive)
4. Verify that the drives on either side of the drive to be replaced have one green LED blinking.
5. Remove the drive to be replaced.
6. Insert the replacement drive (The example replacement drive has UID **ac7c0b2c** as in Figure 121)

7. Take the new drive out of service mode:

   ssau -d a:n -s reset (where **a** is the adapter number and **n** is the SCSI ID of the target drive)

8. Verify that that the green LEDs on the nodes on either side of the replaced drive have stopped blinking.

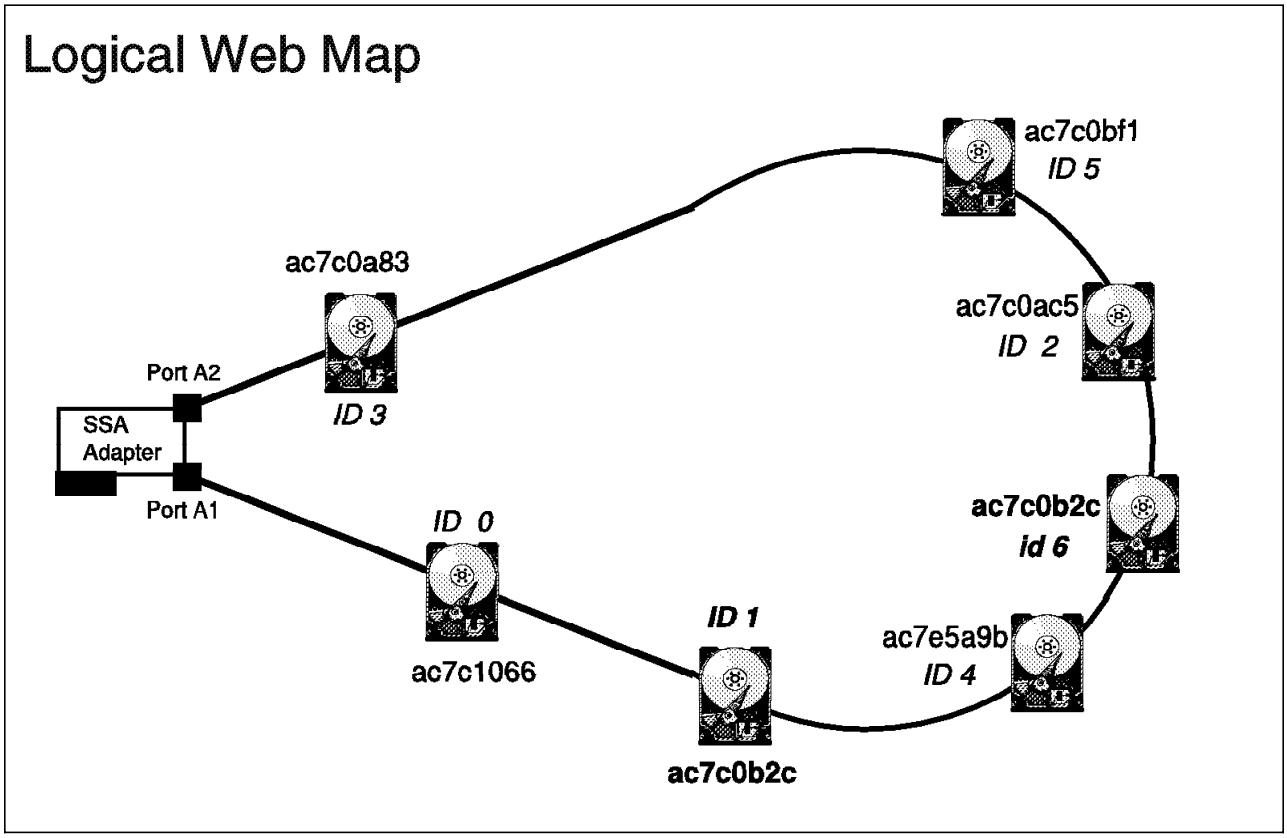

Figure 122. The Result of Replacing a Drive in an SSA Adapter Loop

9. Display the configuration map and compare with Figure 122 to confirm the result:

```
root>
root> sscf -a 0 -d
# IDs of installed adapters
# Configuration map for controller 0
# UID         UID             SCSI
# High   Low             ID
  0004   ac7c1066        00
  0004   ac7c0b2c        01
  0004   ac7c0ac5        02
  0004   ac730a83        03
  0004   ac7b5a9b        04
  0004   ac7c0bf1        05
  0004   ac7c0978        06
root>
```

10. Make the configuration persistent across boots:

    sscf -a 0 -p (for drives connected to adapter 0)

or:

sscf -a 1 -p (for drives connected to adapter 1)

11. Run diagnostics and verify operation of the drive:

    ssau -d a:n -t (where **a** is the adapter number, **n** is the SCSI ID of the drive, and **-t** is the time setting)

12. Attempt a read from the device to make sure it has spun up. You can do this by reading the device to a null device:

```
root>
root> cat /dev/rdsk/c0t0d1s0 > /dev/null
root>
```

for adapter 0

or

```
root>
root> cat /dev/rdsk/c1t0d1s0 > /dev/null
root>
```

for adapter 1.

If the device has not spun up, several **incomplete read- retrying** messages followed by a single **incomplete read- giving up** message will be displayed.

13. The replacement drive is ready to use.

5.2.4.3 Replacing drives in an IBM SSA Interface Controller string.

There are separate procedures to replace drives at the end of a string and to replace drives between nodes on a string. The latter requires you to edit the map file so make sure you have an archived copy of the original.

Replacing a Drive at the End of the Adapter String:

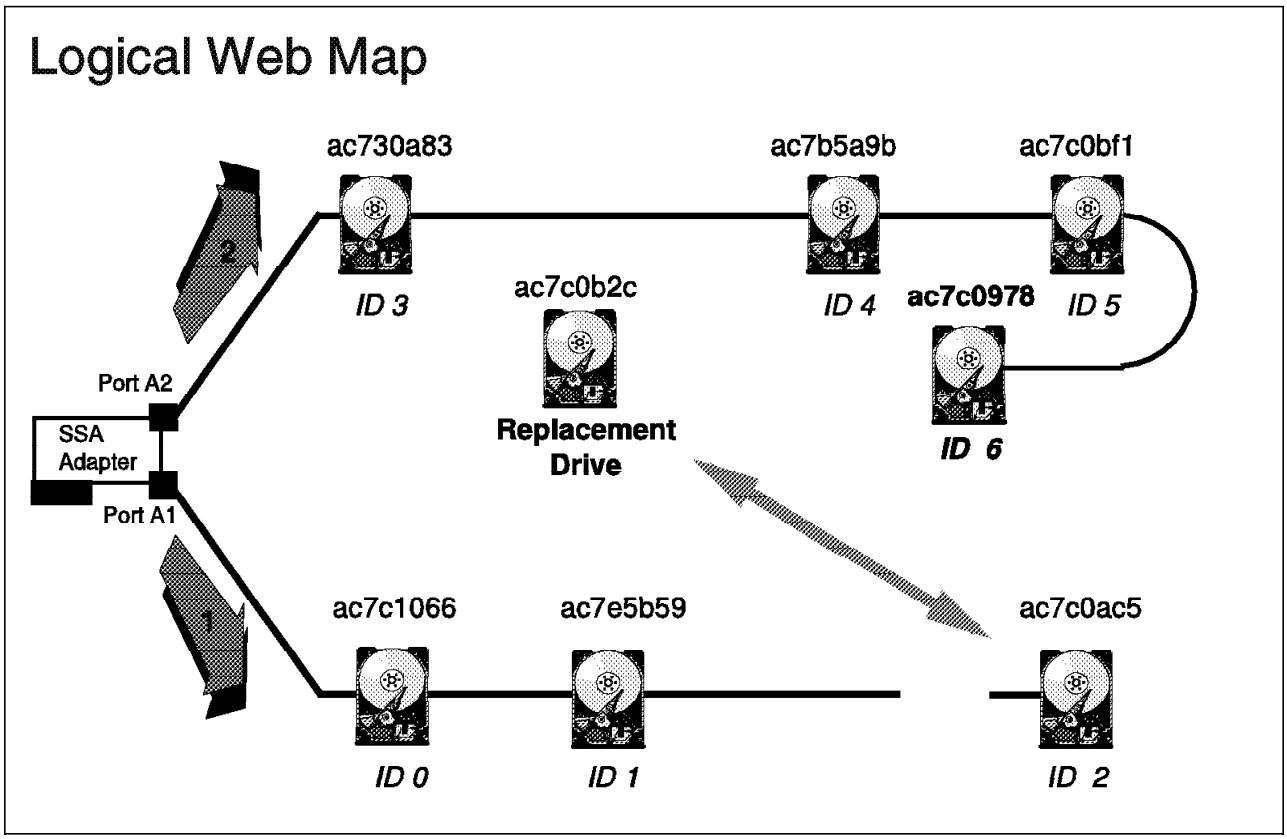

Figure 123. Replacing a drive at the end of the string

Follow this procedure:

1. Identify the drive to be replaced. We have chosen *ID 2* for our example.
2. Reserve the SCSI ID of the drive being removed for the replacement:

   sscf -a 0 -r (for adapter 0)

   or **sscf -a 1 -r** (for adapter 1)
3. Set the drive to be replaced in service mode:

   ssau -d a:n -s set (where **a** is the adapter number and **n** is the SCSI ID of the target drive)
4. Remove the drive to be replaced
5. Attach the new drive
6. Take the new drive out of service mode:

   ssau -d a:n -s reset (where **a** is the adapter number and **n** is the SCSI ID of the target drive)

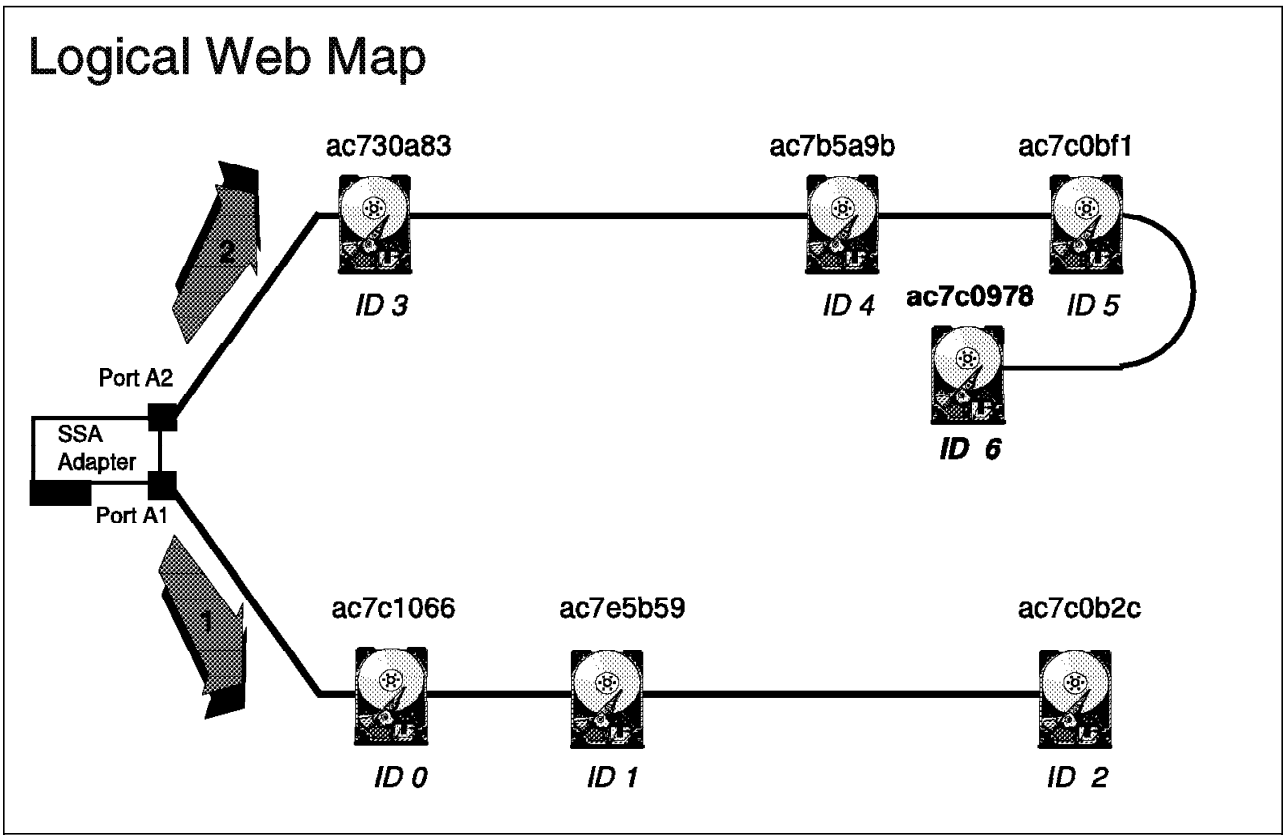

Figure 124. The Resultant Map

Display the configuration map (Figure 124) to confirm the result:

```
root>
root> sscf -a 0 -d
# IDs of installed adapters
# Configuration map for controller 0
# UID         UID             SCSI
# High   Low             ID
  0004   ac7c1066        00
  0004   ac7e5b59        01
  0004   ac7c0b2c        02
  0004   ac730a83        03
  0004   ac7b5a9b        04
  0004   ac7c0bf1        05
  0004   ac7c0978        06
root>
```

7. Make the configuration persistent across boots:

   sscf -a 0 -p (for drives connected to adapter 0)

   or

   sscf -a 1 -p (for drives connected to adapter 1)

8. Run diagnostics and verify operation of the drive:

   ssau -d a:n -t (where **a** is the adapter number

n is the SCSI ID of the drive, and **-t** is the time setting)

9. Attempt a read from the device to make sure it has spun up. You can do this by reading the device to a null device:

```
root>
root> cat /dev/rdsk/c0t0d2s0 > /dev/null
root>
```

for adapter 0

or

```
root>
root> cat /dev/rdsk/c1t0d2s0 > /dev/null
root>
```

for adapter 1

Replacing drives between nodes: This procedure requires that you break the string, thus isolating all the drives downstream from the break. It is therefore important that the system is not active during the procedure.

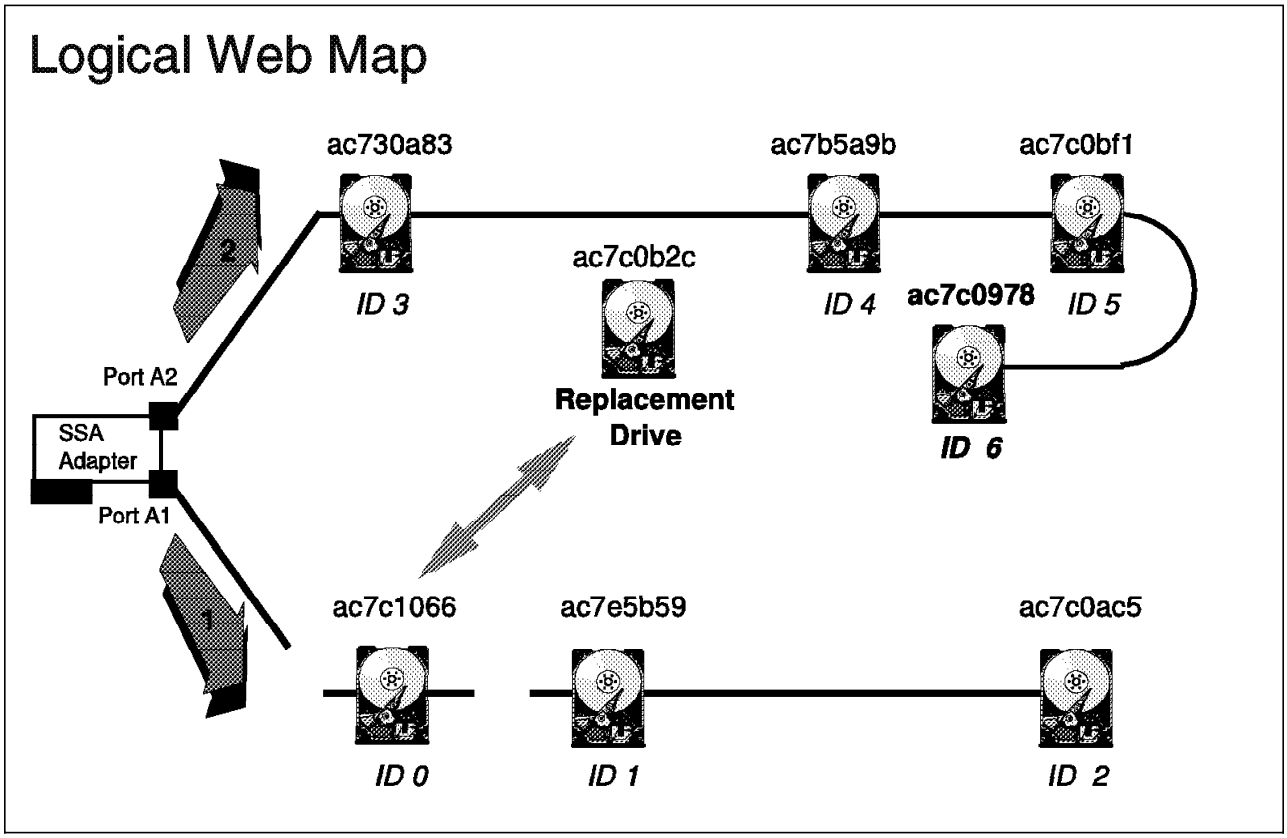

Figure 125. Inserting a Drive Between Nodes on an Adapter String

Use the following procedure:

1. Identify the drive to be replaced. In our example, Figure 119 on page 164, we will replace drive ID 0 with the replacement drive UID **ac7c0b2c** as in Figure 125.
2. Create a map file of the current configuration:

   sscf -a 0 -d > mapfile (for adapter 0)

   or **sscf -a 1 -d > mapfile** (for adapter 1)

   Replace **mapfile** with the file name you choose.
3. Edit the map file with a text editor such as **vi** or **emacs** to replace the UID of the drive you are removing with the drive you are installing.
4. Make the configuration persistent across boots:

   sscf -a 0 -p (for drives connected to adapter 0)

   or

   sscf -a 1 -p (for drives connected to adapter 1)
5. Shut down the Solaris and power down the system.
6. Physically replace the drive, making sure to close the string.
7. Power up the drive and verify that both green LEDs are on and the yellow LED is off.
8. Power up the server and boot the system.

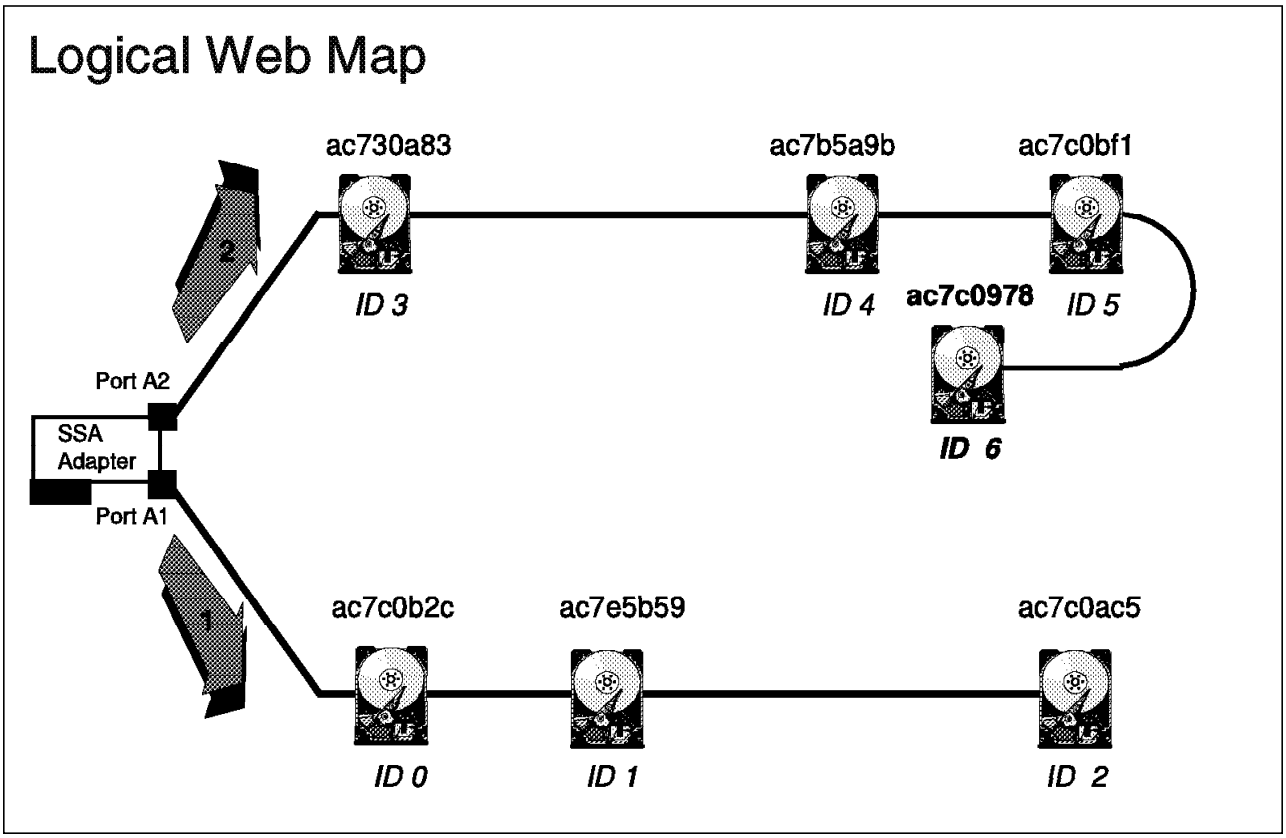

Figure 126. The New Configuration after Replacing ID 0

9. Display the configuration map (Figure 126) to confirm the result:

```
root>
root> sscf -a 0 -d
# IDs of installed adapters
# Configuration map for controller 0
# UID         UID             SCSI
# High   Low             ID
  0004   ac7c0b2c        00
  0004   ac7e5b59        01
  0004   ac7c0ac5        02
  0004   ac730a83        03
  0004   ac7b5a9b        04
  0004   ac7c0bf1        05
  0004   ac7c0978        06
root>
```

10. Run diagnostics and verify operation of the drive:

    ssau -d a:n -t (where **a** is the adapter number **n** is the SCSI ID of the drive, and **-t** is the time setting)

11. Attempt a read from the device to make sure it has spun up. You can do this by reading the device to a null device:

```
root>
root> cat /dev/rdsk/c0t0d2s0 > /dev/null
root>
```

    for adapter 0

    or:

```
root>
root> cat /dev/rdsk/c1t0d2s0 > /dev/null
root>
```

    for adapter 1

12. If the diagnostic procedure is complete, the drive is ready for use.

5.2.5 Loading Firmware

This section describes how to download firmware to the IBM SSA Interface Controller for Sun SBus adapter or to drives attached to it. Ensure that you run firmware downloads from a UPS-assisted system to avoid power cycles during the download. There are separate procedures for downloading disk firmware and downloading adapter firmware.

5.2.5.1 Loading Adapter Firmware

This can be done in two different ways, both achieving the same result. The only time you should download firmware from the Open Boot monitor is when the Solaris operating system is not available.

Note

You must stop all I/O activity directed to the adapter to be updated. We recommend that you temporarily unmount the file systems on the attached disk drives.

Downloading adapter firmware with the* SSAU** utility:*** During the CDROM package installation, the firmware files are copied into the **/opt/PLTsunI/firmware** directory. The command syntax is:

ssau -a n -l (file name) (,f) (where n is the adapter number, **file name** is the path and name of the firmware file, and **,f** is an option to force the update if the firmware being downloaded is older than the existing firmware.

Use the following procedure:

1. Verify the exact path for the firmware image.
2. Stop all SSA activity on the SSA network attached to the IBM SSA Interface Controller for Sun SBus adapter.
3. Type the command from the command shell:

   ssau -a 0 -i (filename).flash (Update the firmware on adapter 0)

   or

   ssau -a 1 -l (filename).flash (force the update of adapter 1 firmware)

   The utility will only download firmware to drives with a firmware that is equal or older in revision. The force option overrides the default and forces the firmware download to all target drives.
4. Execute the adapter diagnostics test to confirm that the adapter is operational. (5.2.6.1, “Adapter Diagnostics” on page 177)
5. Display the SSA web map to confirm that the SSA disk drives are available (“Display the SSA web map (-w)” on page 152)
6. If the diagnostics and the display web map procedures complete without errors, the firmware update has completed successfully.

Downloading adapter firmware from the openboot prompt: You can use the Fcode Flash Update utility to download adapter firmware from the openboot prompt if the Solaris operating system is not available.

Note

You must stop all SSA disk drive activity on the host system, and if the host system is part of a multi-initiator configuration, you must disconnect the SSA cables from the adapter during the download procedure.

The Fcode Flash update utility requires that the update files be accessible from an FTP host, using the TFTP protocol.

Use the following procedure:

1. Copy the update files to the FTP host and verify connectivity between the TFTP host and the local system. The update distribution diskette must

include the file **SSA_UD.PKG** (this is an embedded Forth and Fcode utility file). Other files in the distribution diskette are typically **SSA_FW.UPD, SSA_ST.UPD, SSA_OB.UPD,** and **SSA_ALL.UPD**. These files must be placed in the **/tftpboot** directory of the TFTP host server.

2. From the Forth monitor on the target system, type the following commands:

   ok hex 4000 dload SSA_UD.PKG **ok** 4000 1 byte-load **ok**

3. To update the first file, type the following command:

   ok board bb bus ss update**:**filename Replace bb with the actual board number, replace ss with the actual bus number, and replace filename with the actual file name.

4. The host system will display similar output to the following :

   Sunlight Fcode Update Utility v1.10 Found Adapter on SBus ss of board bb. Downloading file... Updating... Done. ok

5. For each additional file to update on the currently selected Adapter type:

   ok update:filename

   The host will again display text similar to above.

6. Repeat Steps 4-6 for each additional adapter to be updated.

7. Reset and reboot the host system.

8. Test the adapter as follows:

   ok board bb bus ss ssa-test

   The host should reply with appropriate pass, fail, and warning messages.

9. If an error is reported, it may be possible to recover from the current setup by typing :

   ok board bb bus ss ssa-ok?

   The host will indicate a pass/fail status. If a pass is returned, try repeating the flash update process. If a fail is returned, contact your service representative.

5.2.5.2 Downloading Disk Drive Firmware

Only specified levels of disk firmware are supported. Firmware levels are defined in the install package CD- ROM README file. During the CDROM package installation, the firmware files are copied into the **/opt/PLTsunI/firmware** directory. Newer levels of disk firmware may be available. (refer to the IBM Storage Web site)

The syntax of the command is:

ssau -d n -I (filename) (,f)

Replace **n** with the adapter number, replace **(filename)** with the actual file name, and use the **,f** option if you wish to run the procedure in force mode.

Use the following procedure:

1. Verify the exact path of the file that contains the firmware image.
2. Halt all I/O to the drive to be updated. The safest method of achieving this is to temporarily unmount the filesystems on the drive.

3. Type the download command:

   `ssau -d 5 -I dchc85.bin` (update firmware on Disk 5 of Adapter 0)

4. Execute disk diagnostics as in 5.2.6.2, “Disk Drive Diagnostics” to confirm that the drive is operational.
5. Display the Vital Product Data (“Display disk drive Vital Product Data (VPD) (-v)” on page 153) of the drive to confirm that the new firmware is loaded.
6. The procedure is complete.

5.2.6 Troubleshooting

Once you have successfully completed the procedures in Chapter 4, “Configuration,” the system should be operational. This section focuses on post-installation problem determination.

The purpose of the problem determination procedure is to categorize the problem as a host system problem, an SSA storage subsystem problem, or an adapter card problem so you can make the appropriate action and/or service call.

Note

We refer to the *IBM SSA Interface Controller for Sun SBus Installation and Users Guide*, so make sure it is is available.

You may be required to power cycle one or more disk drives so take appropriate measures to ensure that no I/O takes place to the affected disks. The safest method to preserve your data during problem determination is to unmount the related file systems.

5.2.6.1 Adapter Diagnostics

If you are suspicious about the adapter′s health, you can use the SSAU to execute the adapter power-on self-test, thereby avoiding a system restart.The POST flashes the adapter LED in a specific sequence to indicate the adapter′s status. The results of POST can be used to isolate the problem and to produce a solution. (See 5.2.6.3, “ Adapter LED Analysis” on page 178 to translate the LED codes.)

Before running the utility, halt all activity to the adapter and on the host system.

The syntax is:

ssau -a n -t (replace **n** with the adapter number. **t** stands for test.)

5.2.6.2 Disk Drive Diagnostics

If a disk drive′s health is suspect, use the SSAU utility to verify the health of the drive. It causes the drive to execute its built-in diagnostic tests.

Halt all activity to the drive before you execute the command.

The syntax is:

ssau -d a:n -t(run diagnostic test on disk drive n on adapter a)

You can use the results of the test to isolate your problem

5.2.6.3 Adapter LED Analysis

As in Figure 129 on page 182, the first step is to determine whether the POST reported a fatal error. You can do this by observing the two lights positioned between the SSA ports of the IBM SSA Interface Controller.

Refer to the *Installation and User's Guide* for a complete description of the LED sequence at boot up and LED error codes.

Two patterns indicate a fatal error after POST:

- **Yellow ON** and **green OFF or ON**
- **Yellow ON** and **green flashes a patterned sequence**

Green LED error code: An error during the POST or firmware operation causes the green LED to flash a four-digit error code. Each digit is represented by a flash sequence. Flash sequences are grouped in fours and are separated by a pause.

In detail, the blinking pattern occurs as follows:

1. LED is on for 2 seconds
2. LED is off for 1/2 second
3. The LED blinks one to nine times for the first digit. A blink consists of turning the LED on for 1/4 second, and then off for 1/4 second.
4. The LED is off for 1/2 second to separate digits
5. Steps 3 and 4 are repeated for the remaining three digits.

5.2.6.4 Performing a System Health Check

This SSAU function performs a health check on the adapter and its attached drives. The health check consists of the following:

1. A test to check whether the adapter firmware responds to the host driver commands.
2. A request to the adapter firmware for the current web map.
3. A SCSI inquiry command to each disk drive in the SSA web map to determine the redundant power and cooling status of the SSA storage system.

The SSAU health check can be run once, or repeated at a specified interval.

Command line syntax:

ssau -c (interval(repeat count))

If an **interval** is specified, health checks are performed periodically at the interval specified. If the interval is not specified, a single set of health checks are performed.

If an interval is specified, the **repeat count** specifies the number of health checks that are performed.

If an interval is specified without a repeat count, the health checks are performed indefinitely at the specified interval.

The Solaris system log. If any of the above checks fails to complete or returns a bad status, the health check updates the Solaris system log /var/adm/messages.

5.2.6.5 Collecting System Information

When you call a service representative, you will be required to provide certain information before your problem can be addresed. We advise that you collect this information before you start your own problem-solving procedure. The information you record will be used in the problem determination procedure that we provide.

The following is a basic template of the information that you should document:

- Customer name
- Contact person and telephone number
- Host type and model number
- Operating system type and revision level
- SCSI adapter model numbers
- IBM SSA Interface Controller for Sun SBus serial number
- IBM SSA Interface Controller for Sun SBus vital product data (VPD)
- SSA Storage Subsystem model number and serial number
- Description and configuration of SSA topology information such as SSA loop structure and number of drives (see “Display the SSA web map (-w)” on page 152)
- Record file subsystem information such as mirroring and other raid levels
- Description of problem
- LED POST pattern (5.2.6.3, “ Adapter LED Analysis” on page 178)
- If an individual drive was identified as defective by a message on the host monitor or a recorded SCSI error status in the Operating System log, record the failed SCSI/LUN as seen from the host

5.2.6.6 The Troubleshooting Procedure

The purpose of this procedure is to isolate the problem so that the appropriate action can be taken to correct it. We achieve this by obtaining an SRN (If the current state of the SSA system allows us to) that identifies either a disk problem, a subsystem problem, or an adapter problem.

We have divided the problem determination procedure into three sections. The main procedure is the backbone of the problem determination procedure. If you follow the steps in the main procedure, you will be pointed to a second procedure, depending on whether or not you have identified a fatal error. During one cycle of the problem determination procedure, you will use only two out of three of the following sections:

The main problem determination procedure:

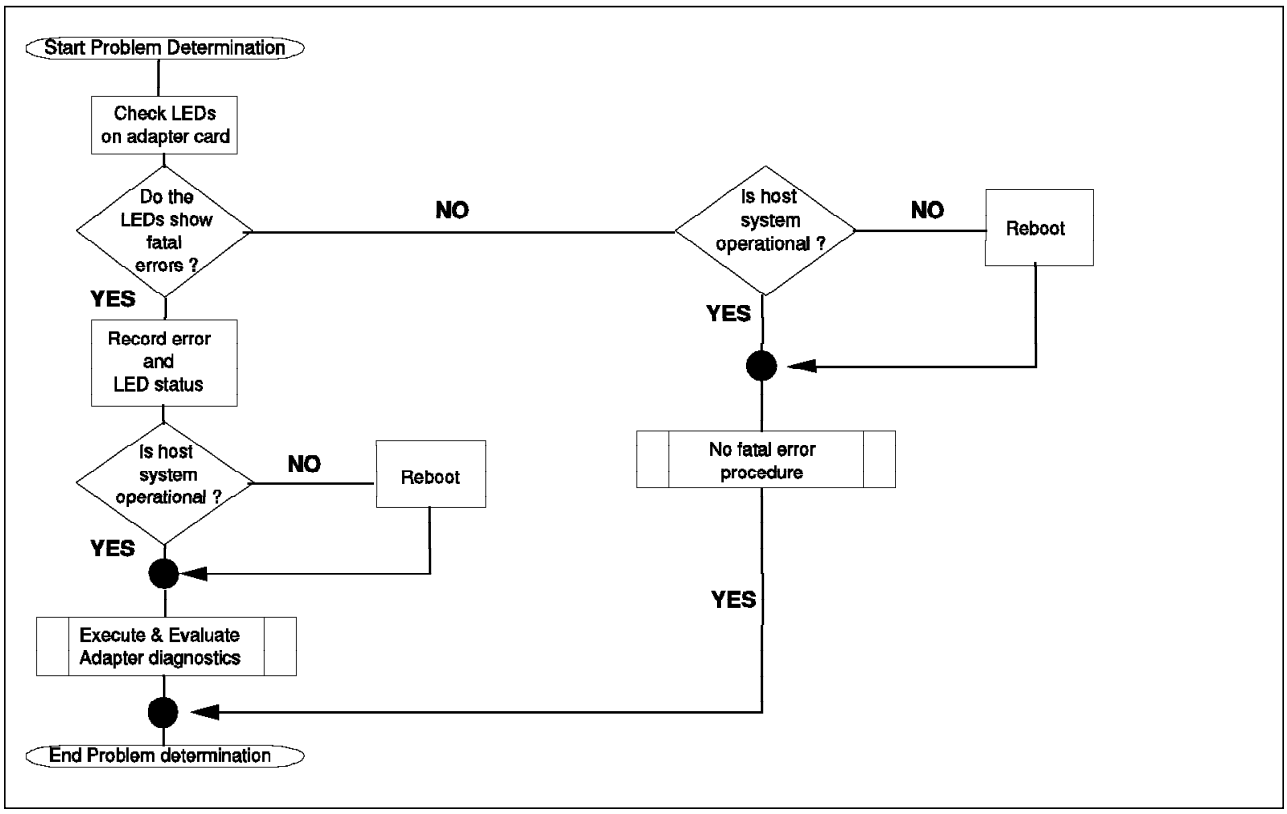

Figure 127. The Main Problem Determination Procedure

As you can see in Figure 127, the purpose of running the main problem determination procedure is to verify whether there is a fatal adapter error or not and to run the appropriate procedure to correct the problem.

The following steps describe the flow chart in a hierarchical order:

- **Start Problem Determination**. Make sure that you have sufficient tools to document your findings and that you have the *IBM SSA Interface Controller for Sun SBus Installation and User's Guide* available.
- **Check the LEDs on the adapter card**. Record the LED sequence at boot up and check the LED status according to 5.2.6.3, " Adapter LED Analysis" on page 178 in order to confirm whether or not there is a fatal adapter error.
- **Do the LEDs show fatal errors?** If no fatal error is reported and the host system is operational, run the 'No fatal error procedure.' If no fatal error is reported but the host system is not operational, reboot and then run the 'No fatal error procedure.'
- **Record error and LED status**. You need to have this information to continue the problem determination procedure. The LED status of the drives, power supplies, power cards, and bypass cards should be checked.
- **Is host system operational?** If the host system is not operational, you need to boot it into a usable state. Reboot the system to correct the problem. If the problem is still not corrected, contact your service host system service representative.

  If the system is operational, proceed to the Fatal error procedure to diagnose and correct the fatal error.

- **End Problem determination**. If your SSA network is still problematic, record the system information as described in 5.2.6.5, “Collecting System Information” on page 179 and call your service representative.

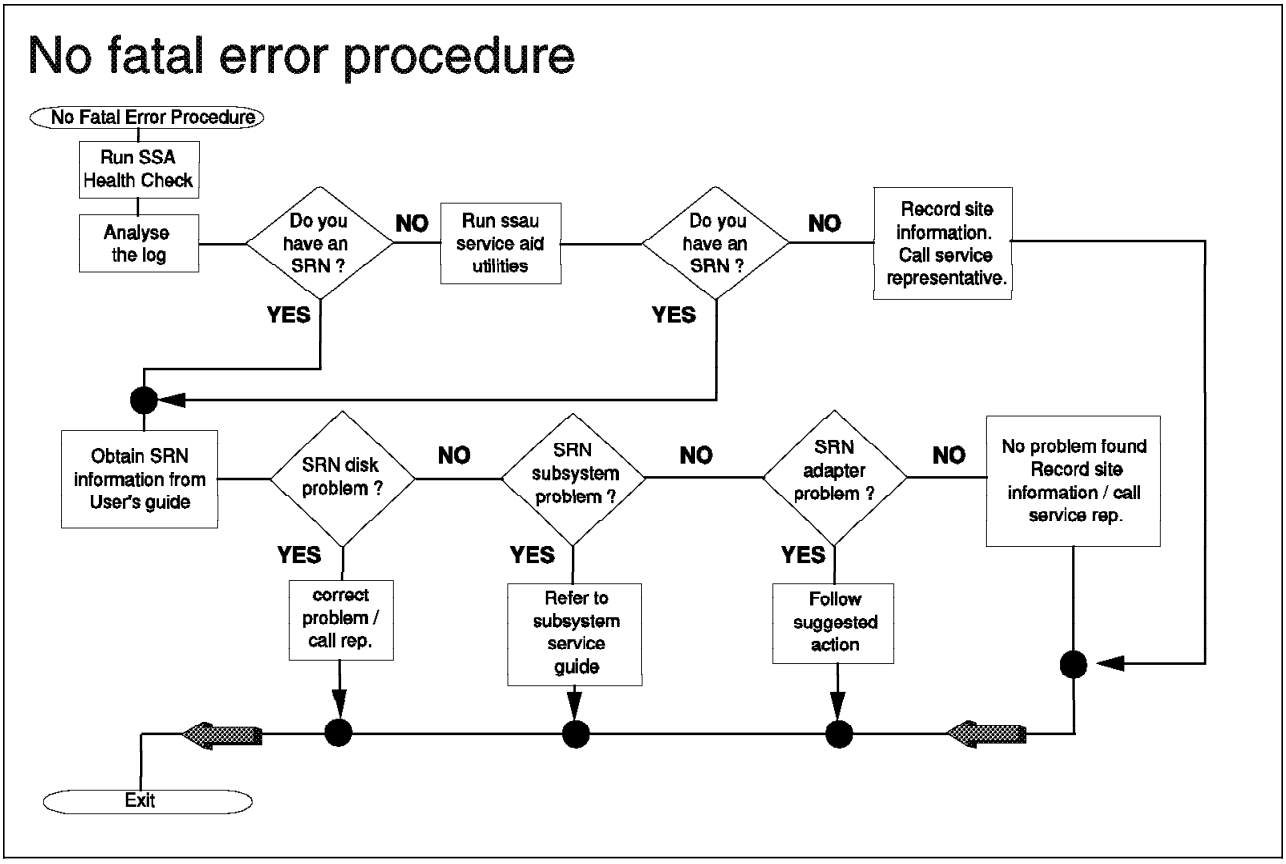

Figure 128. Run this Procedure if You Have Not Had a Fatal Adapter Error

No fatal error procedure: The flow chart, Figure 128, is a representation of the procedure you take to associate the problem with a specific device.

The following steps describe the flowchart in a hierarchical order:

- **Run the SSA health check**. Run the procedure described in 5.2.6.4, “Performing a System Health Check” on page 178, to test the adapter and to log all errors in the Solaris system message log.
- **Analyze the log**. Analyze the Solaris system message log to check whether a SRN has been generated. If no SRN is reported, run the SSAU utility on all the suspected device and analyze the Solaris system message log again. If there are still no SRNs reported, record the system information as in 5.2.6.5, “Collecting System Information” on page 179. If an SRN is reported, continue the procedure.
- **Obtain SRN information from User′s guide**. Use the guide to determine what part of the system is causing the problem and correct the problem or call the appropriate service representative.

No fatal error procedure t oidentify an adapter problem: Use this procedure to run diagnostics on the adapter to determine the corrective action.

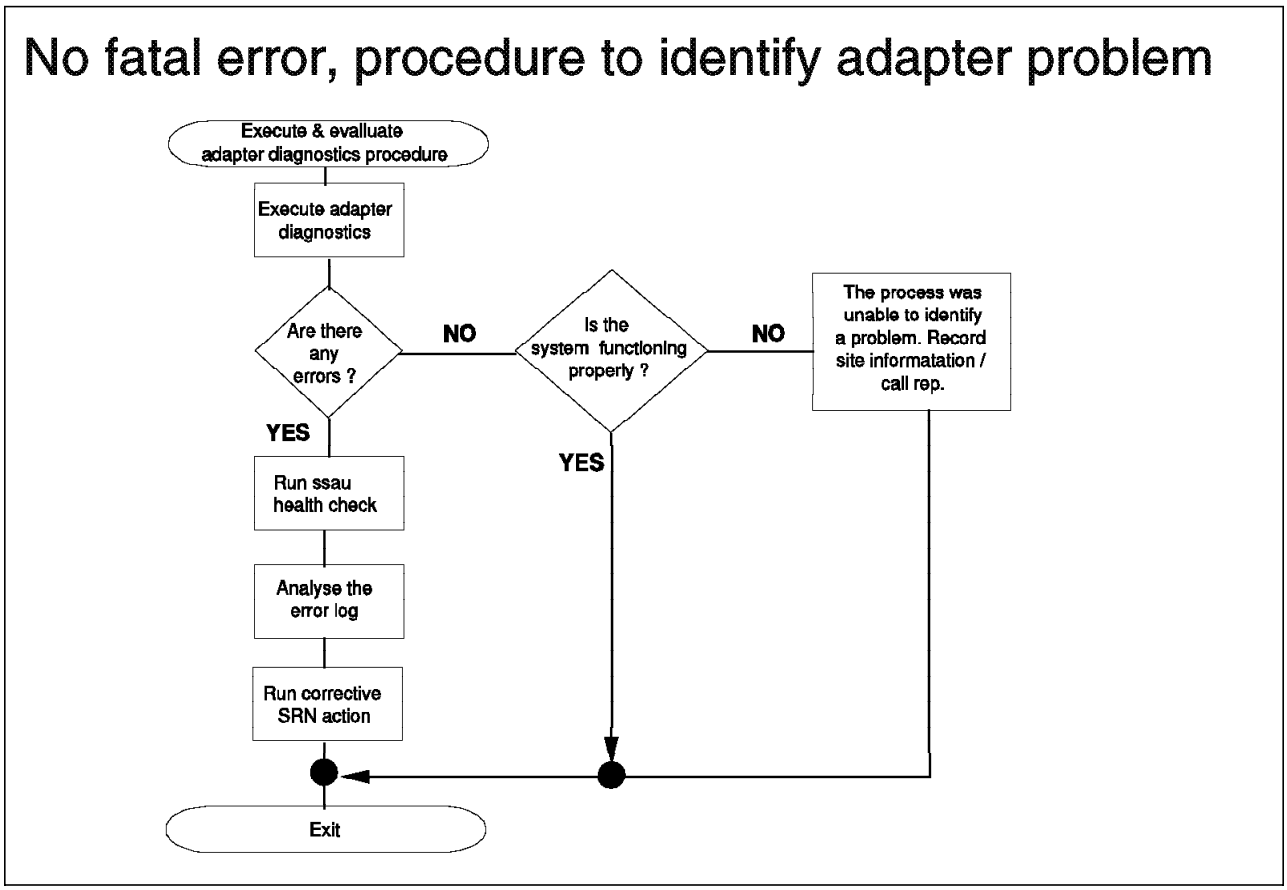

Figure 129. No Fatal Error, Procedure to Identify Adapter Problem

This procedure is based upon the assumption that the problem is adapter specific. The following steps describe the flow chart in a hierarchical order:

- **Execute adapter diagnostics**. Use the procedure in 5.2.6.1, “Adapter Diagnostics” on page 177, to execute the adapter POST. If the POST does not report any errors, check again whether the system is functioning. If the system still isn’t functioning after POST runs without errors, record the system information (5.2.6.5, “Collecting System Information” on page 179) and call your service representative. If an SRN/s was reported, continue the procedure.
- **Run SSAU health check**. Run the procedure described in 5.2.6.4, “Performing a System Health Check” on page 178, and analyze the Solaris system message log.
- **Run corrective SRN action**. Refer to your *Installation and User’s Guide*.

Chapter 6. Migration and Coexistence Guidelines

In this chapter we explore migrating

- from 7190 Model 100 to SSA Interface Controller for SBus
- from SCSF to 7190 Model 100 and SSA Interface Controller for SBus
- from SSA Interface Controller for SBus to 7190 Model 100

6.1 Migration from SCSI to 7190 Model 100

To migrate data from a SCSI environment to the 7190 Model 100 environment, the guidelines listed below should be helpful. Along with the guidelines, we provide examples to illustrate the steps:

1. Back up all of your data.

   We believe data integrity to be of the utmost importance. Although we would not expect problems in implementing this procedure, you should always have a fallback option. Having a good backup of your data gives you this option.

2. Using the **df -ak command**, create a log of all file systems associated with the drives to be replaced. Identify the device path names and their associated mount points. Figure 130 shows the output of the **df -ak** command.

df -ak Command (SCSI to 7190)

```
# df -ak > df.data
# view df.data
 Filesystem            kbytes    used     avail   capacity   Mounted on
 /dev/dsk/c0t4d0s0     288523    15534    244139     6%      /
 /dev/dsk/c0t4d0s6     384883    81700    264703     24%     /usr
 /proc                 0         0        0          0%      /proc
 fd                    0         0        0          0%      /dev/fd
 /dev/dsk/c0t4d0s4     481243    14604    418519     4%      /var
 /dev/dsk/c0t4d0s5     404363    47419    316514     14%     /opt
 /dev/dsk/c0t4d0s3     384883    106488   239915     31%     /usr/openwin
 swap                  211520    36       211484     1%      /tmp
 /dev/dsk/c2t0d0s6     8365850   9        7529261    1%      /filesys1
 /dev/dsk/c2t1d0s6     4139980   9        3725981    1%      /filesys2
 /dev/dsk/c2t2d0s6     4139980   9        3725981    1%      /filesys3
 /dev/dsk/c2t3d0s6     4139980   9        3725981    1%      /filesys4
 /dev/dsk/c2t4d0s6     4139980   9        3725981    1%      /filesys5
 -hosts                0         0        0          0%      /net
 auto_home             0         0        0          0%      /home
 -xfn                  0         0        0          0%      /xfn
 sc2000:vold(pid232)   0         0        0          0%      /vol
```

Figure 130. Output of

3. If you have made entries in /etc/vfstab to automatically mount the file systems associated with the drives to be replaced, comment these entries out. Figure 131 on page 184 shows an example with the appropriate entries commented out.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
#/dev/dsk/c2t0d0s6  /dev/rdsk/c2t0d0s6  /filesys1  ufs  1  yes  -
#/dev/dsk/c2t1d0s6  /dev/rdsk/c2t1d0s6  /filesys2  ufs  1  yes  -
#/dev/dsk/c2t2d0s6  /dev/rdsk/c2t2d0s6  /filesys3  ufs  1  yes  -
#/dev/dsk/c2t3d0s6  /dev/rdsk/c2t3d0s6  /filesys4  ufs  1  yes  -
#/dev/dsk/c2t4d0s6  /dev/rdsk/c2t4d0s6  /filesys5  ufs  1  yes  -
```

Figure 131. Sample /etc/vfstab File (SCSI to 7190)

4. Shut down the system. Install the 7190 Model 100 adapter in accordance with the *7190 Model 100 SCSI Host to SSA Loop Attachment Model 100 Installation and User's Guide*. SSA subsystems should be configured and attached in the desired configuration.
5. Perform the Sun host steps outlined in the "Connecting to a SUN Host System" section of the *7190 Model 100 SCSI Host to SSA Loop Attachment Model 100 Installation and User's Guide*. In particular:
   a. Identify the SCSI ID/LUN pairs mapped by the 7190.

      Issue the Forth command **probe-scsi-all** at the **ok** prompt.

   b. Update the sd.conf file with the appropriate SCSI ID/LUN pairs.
   c. Create at least one 7190 Model 100 partition for each 7190 Model 100 on the system.
   d. Create a 7190 Model 100 configuration file, 7190.cfg, for each 7190.
   e. Install the service functions application.
6. After the service functions application has been installed, use the format command to confirm the new SCSI ID/LUN pairs assigned to the SSA drives. These should match the pairs observed from the **probe-scsi-all** command.
7. Using the output of the **df -ak** command from Step 2 on page 183, along with the output of the format command from Step 6, create a table to show the relationship between the original SCSI device path names and the new 7190 Model 100 path names for the SSA devices. Include the applicable mount points. An example is shown in Figure 132 on page 185.

```
Original SCSI      Mount        7190 model 100 (New)
Device Names       Points       Device Names
-----------------  -----------  --------------

/dev/dsk/c2t0d0s6  /filesys1    /dev/dsk/c4t0d0s6
/dev/dsk/c2t1d0s6  /filesys2    /dev/dsk/c4t0d1s6
/dev/dsk/c2t2d0s6  /filesys3    /dev/dsk/c4t0d2s6
/dev/dsk/c2t3d0s6  /filesys4    /dev/dsk/c4t0d3s6
/dev/dsk/c2t4d0s6  /filesys5    /dev/dsk/c4t0d4s6
```

Figure 132. Original SCSI and New 7190 Model 100 Device Names.

8. Using the table you have created, update all applications that use a device path name to reference the respective partitions or mount points. As a minimum, you have to update /etc/vfstab if you are using it to automatically mount partitions. Figure Figure 133 shows an example of the new /etc/vfstab.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
/dev/dsk/c4t0d0s6  /dev/rdsk/c4t0d0s6  /filesys1  ufs  1  yes  -
/dev/dsk/c4t0d1s6  /dev/rdsk/c4t0d1s6  /filesys2  ufs  1  yes  -
/dev/dsk/c4t0d2s6  /dev/rdsk/c4t0d2s6  /filesys3  ufs  1  yes  -
/dev/dsk/c4t0d3s6  /dev/rdsk/c4t0d3s6  /filesys4  ufs  1  yes  -
/dev/dsk/c4t0d4s6  /dev/rdsk/c4t0d4s6  /filesys5  ufs  1  yes  -
```

Figure 133. Updated /etc/vfstab for SCSI to 7190 Model 100

9. Incorporate the updated device names into the kernel configuration. There are two ways to accomplish this. One is to issue this command:

   reboot -- -r

   The other is to issue the following series of commands:

   drvconfig

   devlinks

   disks

   ucblinks

   The method you use depends on your particular system environment.

6.2 Migration from SCSI to IBM SSA Interface Controller

To migrate data from a SCSI environment to the IBM SSA Interface Controller for Sun SBus environment, the guidelines listed below should be helpful. Along with the guidelines, we provide examples to illustrate the steps.

1. Back up all of your data.

   Once again, we believe data integrity to be of the utmost importance. Although we would not expect problems in implementing this procedure, you should always have a fallback option. Having a good backup of your data gives you this option.

2. Using the **df -ak** command, create a log of all file systems associated with the drives to be replaced. Identify the device path names and their associated mount points. Figure 134 shows the output of the **df -ak** command.

df -ak Command (SCSI to SSA Interface Controller)

```
# df -ak > df.data
# view df.data
 Filesystem            kbytes    used     avail   capacity   Mounted on
 /dev/dsk/c0t4d0s0     288523    15534    244139     6%      /
 /dev/dsk/c0t4d0s6     384883    81700    264703     24%     /usr
 /proc                 0         0        0          0%      /proc
 fd                    0         0        0          0%      /dev/fd
 /dev/dsk/c0t4d0s4     481243    14604    418519     4%      /var
 /dev/dsk/c0t4d0s5     404363    47419    316514     14%     /opt
 /dev/dsk/c0t4d0s3     384883    106488   239915     31%     /usr/openwin
 swap                  211520    36       211484     1%      /tmp
 /dev/dsk/c4t0d0s6     8365850   9        7529261    1%      /filesys1
 /dev/dsk/c4t0d1s6     4139980   9        3725981    1%      /filesys2
 /dev/dsk/c4t0d2s6     4139980   9        3725981    1%      /filesys3
 /dev/dsk/c4t0d3s6     4139980   9        3725981    1%      /filesys4
 /dev/dsk/c4t0d4s6     4139980   9        3725981    1%      /filesys5
 -hosts                0         0        0          0%      /net
 auto_home             0         0        0          0%      /home
 -xfn                  0         0        0          0%      /xfn
 sc2000:vold(pid232)   0         0        0          0%      /vol
```

Figure 134. Output of

3. If you have made entries in /etc/vfstab to automatically mount the file systems associated with the drives to be replaced, comment these entries out. Figure 135 on page 187 shows an example with the appropriate entries commented out.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
#/dev/dsk/c4t0d0s6  /dev/rdsk/c4t0d0s6  /filesys1  ufs  1  yes  -
#/dev/dsk/c4t0d1s6  /dev/rdsk/c4t0d1s6  /filesys2  ufs  1  yes  -
#/dev/dsk/c4t0d2s6  /dev/rdsk/c4t0d2s6  /filesys3  ufs  1  yes  -
#/dev/dsk/c4t0d3s6  /dev/rdsk/c4t0d3s6  /filesys4  ufs  1  yes  -
#/dev/dsk/c4t0d4s6  /dev/rdsk/c4t0d4s6  /filesys5  ufs  1  yes  -
```

Figure 135. Sample /etc/vfstab File (SCSI to SSA Interface Controller)

4. Shut down the system. Install the IBM SSA Interface Controller for Sun SBus adapter in accordance with the *IBM SSA Interface Controller for Sun SBus Installation and User's Guide*.
5. Boot the system and install the SBus driver and utilities package in accordance with the *IBM SSA Interface Controller for Sun SBus Installation and User's Guide*.
6. After the package has been installed and the system rebooted, use the format command to confirm the new SCSI ID/LUN pairs assigned to the SSA drives.
7. Create a new map file for the drives, using the **sscf -d** command. Figure 136 demonstrates this command and its output.

```
# sscf -a 0 -d > mapfile.0
# view mapfile.0
 # IDs of installed Adapters: 0 1
 # Configuration map for Adapter 0
 # UID     UID    SCSI
 # High    Low     ID
   0004  ac7e4fab  00
   0004  ac7e50cb  01
   0004  ac7e4f49  02
   0004  ac7e50a8  03
   0004  ac9de6d8  04
```

Figure 136. Using SSCF to Create a Map File (SCSI to SSA Interface Controller)

8. If you want to change the SCSI target IDs assigned to the new drives, edit the map file created by the **sscf** command. Assign the desired SCSI target IDs to the respective SSA drives. Figure 137 on page 188 shows the previous map file edited to reflect a rearranged configuration from the original IDs.

   Note that the IBM SSA Interface Controller for Sun SBus maps drives to LUN 0 on SCSI target IDs 0 through 47. If this changes the SCSI controller/ID/LUN combinations from the original SCSI configuration (which it probably will),

and your system has any applications that make direct access to device path names, you must update the application references to reflect the updated path names. This task can vary in complexity depending on your particular applications.

```
# IDs of installed Adapters: 0 1
# Configuration map for Adapter 0
# UID     UID    SCSI
# High    Low     ID
  0004  ac9de6d8  00
  0004  ac7e50a8  01
  0004  ac7e4f49  02
  0004  ac7e50cb  03
  0004  ac7e4fab  04
```

Figure 137. Updating SBus Map File to New Target ID Assignments

9. Use the **sscf** utility with the -p (permanent) and -m (mapfile) options to make your edited map file permanent. The format of this command is:

   sscf -a adapter_no -p -m mapfile_name

   For example, **sscf -a 0 -p -m mapfile.0**

10. Using the output of the **df** command from Step 2 on page 186, along with the updated map file, create a table to show the relationship between the original SCSI device path names and the new SBus adapter path names for the SSA devices. Include the applicable mount points. An example is shown in Figure 138.

```
Original SCSI      Mount        SBus Interface
Device Names       Points       Device Names
-----------------  -----------  ---------------

/dev/dsk/c4t0d0s6  /filesys1    /dev/dsk/c2t0d0s6
/dev/dsk/c4t0d1s6  /filesys2    /dev/dsk/c2t1d0s6
/dev/dsk/c4t0d2s6  /filesys3    /dev/dsk/c2t2d0s6
/dev/dsk/c4t0d3s6  /filesys4    /dev/dsk/c2t3d0s6
/dev/dsk/c4t0d4s6  /filesys5    /dev/dsk/c2t4d0s6
```

Figure 138. Original SCSI and New SBus Adapter Device Names

11. Using the table you have created, update all applications that use a device path name to reference the respective partitions or mount points. As a minimum, you will have to update /etc/vfstab if you are using it to automatically mount partitions. Figure 139 on page 189 shows an example of the new /etc/vfstab for our sample test environment.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
/dev/dsk/c2t0d0s6  /dev/rdsk/c2t0d0s6  /filesys1  ufs  1  yes  -
/dev/dsk/c2t1d0s6  /dev/rdsk/c2t1d0s6  /filesys2  ufs  1  yes  -
/dev/dsk/c2t2d0s6  /dev/rdsk/c2t2d0s6  /filesys3  ufs  1  yes  -
/dev/dsk/c2t3d0s6  /dev/rdsk/c2t3d0s6  /filesys4  ufs  1  yes  -
/dev/dsk/c2t4d0s6  /dev/rdsk/c2t4d0s6  /filesys5  ufs  1  yes  -
```

Figure 139. Updated /etc/vfstab for SCSI to SBus Adapter.

12. Incorporate the updated device names into the kernel configuration. There are two ways to accomplish this. One is to issue this command:

    reboot -- -r

    The other is to issue the following series of commands:

    drvconfig

    devlinks

    disks

    ucblinks

    The method you use depends on your particular system environment.

6.3 Migration from SSA Interface Controller to 7190 model 100

To migrate data from the IBM SSA Interface Controller for Sun SBus environment to the 7190 Model 100 environment, the guidelines listed below should be helpful. Along with the guidelines, we provide examples of a sample test environment.

1. Back up all your data.

   As we have stated before, we believe data integrity to be of the utmost importance. Although we would not expect problems in implementing this procedure, you should always have a fallback option. Having a good backup of your data gives you this option.

2. Make a copy of the SBus adapter topology, using the **sscf -d** command. Figure 140 on page 190 demonstrates this command and its output. Do this for all IBM SSA Interface Controller for Sun SBus adapters in the system.

sscf to Create a Map File (SBus adapter to 7190)

```
# sscf -a 0 -d > mapfile.0
# view mapfile.0
 # IDs of installed Adapters: 0 1
 # Configuration map for Adapter 0
 # UID     UID    SCSI
 # High    Low     ID
   0004  ac7e4fab  00
   0004  ac7e50cb  01
   0004  ac7e4f49  02
   0004  ac7e50a8  03
   0004  ac9de6d8  04
```

Figure 140. Using

3. Using the **df -ak** command, create a log of all file systems associated with SSA drives. Identify the device path names and their associated mount points. An example of this output is shown in Figure 141.

df -ak Command (SBus Adapter to 7190)

```
# df -ak > df.data
# view df.data
 Filesystem            kbytes    used     avail   capacity   Mounted on
 /dev/dsk/c0t4d0s0     288523    15534    244139     6%      /
 /dev/dsk/c0t4d0s6     384883    81700    264703     24%     /usr
 /proc                 0         0        0          0%      /proc
 fd                    0         0        0          0%      /dev/fd
 /dev/dsk/c0t4d0s4     481243    14604    418519     4%      /var
 /dev/dsk/c0t4d0s5     404363    47419    316514     14%     /opt
 /dev/dsk/c0t4d0s3     384883    106488   239915     31%     /usr/openwin
 swap                  211520    36       211484     1%      /tmp
 /dev/dsk/c2t0d0s6     8365850   9        7529261    1%      /filesys1
 /dev/dsk/c2t1d0s6     4139980   9        3725981    1%      /filesys2
 /dev/dsk/c2t2d0s6     4139980   9        3725981    1%      /filesys3
 /dev/dsk/c2t3d0s6     4139980   9        3725981    1%      /filesys4
 /dev/dsk/c2t4d0s6     4139980   9        3725981    1%      /filesys5
 -hosts                0         0        0          0%      /net
 auto_home             0         0        0          0%      /home
 -xfn                  0         0        0          0%      /xfn
 sc2000:vold(pid232)   0         0        0          0%      /vol
```

Figure 141. Output of

4. If you have made entries in /etc/vfstab to automatically mount the file systems associated with the SSA drives, comment these entries out. Figure 142 on page 191 shows our example with the appropriate entries commented out.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
#/dev/dsk/c2t0d0s6  /dev/rdsk/c2t0d0s6  /filesys1  ufs  1  yes  -
#/dev/dsk/c2t1d0s6  /dev/rdsk/c2t1d0s6  /filesys2  ufs  1  yes  -
#/dev/dsk/c2t2d0s6  /dev/rdsk/c2t2d0s6  /filesys3  ufs  1  yes  -
#/dev/dsk/c2t3d0s6  /dev/rdsk/c2t3d0s6  /filesys4  ufs  1  yes  -
#/dev/dsk/c2t4d0s6  /dev/rdsk/c2t4d0s6  /filesys5  ufs  1  yes  -
```

Figure 142. Sample /etc/vfstab File (SBus Adapter to 7190)

5. Shut down the system. Install the 7190 Model 100 adapter in accordance with the *7190 Model 100 SCSI Host to SSA Loop Attachment Model 100 Installation and User's Guide*. SSA subsystems should be configured and attached in the desired configuration.

   Note that the 7190 Model 100 maps drives to LUNs 0 through 7 on SCSI target IDs 0 through 5, 6 (in mode 0), and 8 through 15. The IBM SSA Interface Controller for Sun SBus maps drives to LUN 0 on SCSI target IDs 0 through 47. Unlike the IBM SSA Interface Controller for Sun SBus, there is no way to reorder drive sequence on the 7190 Model 100 other than by physical arrangement. The only option is to choose between LUN priority and ID priority. Based on this, consider whether a rearrangement of drives within your storage subsystem would be appropriate for your particular application. For further information about 7190 Model 100 drive mapping, refer to Chapter 2, "SBus Adapter and 7190 Technology" on page 9.

6. Perform the Sun host steps outlined in the "Connecting to a SUN host system" section of the *7190 Model 100 SCSI Host to SSA Loop Attachment Model 100 Installation and User's Guide*. In particular:

   a. Identify the SCSI ID/LUN pairs mapped by the 7190.

      Issue the Forth command **probe-scsi-all** at the **ok** prompt.

   b. Update the sd.conf file with the appropriate SCSI ID/LUN pairs.

   c. Create at least one 7190 Model 100 partition for each 7190 Model 100 on the system.

   d. Create a 7190 Model 100 configuration file, 7190.cfg, for each 7190.

   e. Install the service functions application.

7. After the service functions application has been installed, use the format command to confirm the new SCSI ID/LUN pairs assigned to the SSA drives. These should match the combinations observed from the **probe- scsi-all** command. In the sample case we are following, the device path names have changed from the SSA Interface Controller nomenclature of c2t*d0s* to the 7190 Model 100 nomenclature of c4t*d*s*.

8. Using the output of the **df** command from Step 3 on page 190, along with the output of the format command from Step 7, create a table to show the relationship between the original SSA Interface Controller device path names and the new 7190 Model 100 path names for the SSA devices. Include the applicable mount points. A sample for our test environment is shown in Figure 143.

```
SBus (Original)    Mount        7190 model 100 (New)
Device Names       Points       Device Names
-----------------  -----------  ---------------

/dev/dsk/c2t0d0s6  /filesys1    /dev/dsk/c4t0d0s6
/dev/dsk/c2t1d0s6  /filesys2    /dev/dsk/c4t0d1s6
/dev/dsk/c2t2d0s6  /filesys3    /dev/dsk/c4t0d2s6
/dev/dsk/c2t3d0s6  /filesys4    /dev/dsk/c4t0d3s6
/dev/dsk/c2t4d0s6  /filesys5    /dev/dsk/c4t0d4s6
```

Figure 143. IBM SSA Interface Controller and New 7190 model 100 Device Names.

9. Using the table you have created, update all applications that use a device path name to reference the respective partitions or mount points. As a minimum, you have to update /etc/vfstab if you are using it to automatically mount partitions. Figure 144 shows an example of the new /etc/vfstab for our sample test environment.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
/dev/dsk/c4t0d0s6  /dev/rdsk/c4t0d0s6  /filesys1  ufs  1  yes  -
/dev/dsk/c4t0d1s6  /dev/rdsk/c4t0d1s6  /filesys2  ufs  1  yes  -
/dev/dsk/c4t0d2s6  /dev/rdsk/c4t0d2s6  /filesys3  ufs  1  yes  -
/dev/dsk/c4t0d3s6  /dev/rdsk/c4t0d3s6  /filesys4  ufs  1  yes  -
/dev/dsk/c4t0d4s6  /dev/rdsk/c4t0d4s6  /filesys5  ufs  1  yes  -
```

Figure 144. Updated /etc/vfstab for SBus Adapter to 7190 Model 100

10. Incorporate the updated device names into the kernel configuration. There are two ways to accomplish this. One is to issue this command:

    reboot -- -r

    The other is to issue the following series of commands:

    drvconfig

    devlinks

    disks

    ucblinks

The method you use depends on your particular system environment.

6.4 Migration from 7190 model 100 to SSA Interface Controller

To migrate data from the 7190 Model 100 environment to the IBM SSA Interface Controller for Sun SBus environment, the guidelines listed below should be helpful. Along with the guidelines, we provide examples from our sample test environment.

1. Back up all of your data.

   We reiterate that we believe data integrity to be of the utmost importance. Although we would not expect problems in implementing this procedure, you should always have a fallback option. Having a good backup of your data gives you this option.

2. Make a copy of the 7190 Model 100 SSA topology.

   Run the 7190 Model 100 service functions utility, SSA_SERV, and execute the **t** option (show SSA topology). Copy the output of this command and paste it into a map file that will be used later in this process. Figure 145 shows an example map file. Do this for all 7190s in the system.

```
*** SSA TOPOLOGY ***
CTR1.7190.22001069.SM        <--- 7190 model 100 Controller
SSA.PORT.A1=
=N.TOD0.AC9DE6D8.OK=         <-+
=N.TOD1.AC7E50A8.OK=           |
=N.TOD2.AC7E4F49.OK=           |-- SSA Drives
=N.TOD3.AC7E50CB.OK=           |
=N.TOD4.AC7E4FAB.OK=         <-+
=SSA.PORT.A2
```

Figure 145. Sample 7190 Model 100 Topology Map. This is the output of SSA_SERV Option **t**

3. Using the **df -ak** command, create a log of all file systems associated with the SSA drives. Identify the device path names and their associated mount points. An example of this output is shown in Figure 146 on page 194.

```
# df -ak > df.data
# view df.data
 Filesystem            kbytes    used     avail   capacity   Mounted on
 /dev/dsk/c0t4d0s0     288523    15534    244139     6%      /
 /dev/dsk/c0t4d0s6     384883    81700    264703     24%     /usr
 /proc                 0         0        0          0%      /proc
 fd                    0         0        0          0%      /dev/fd
 /dev/dsk/c0t4d0s4     481243    14604    418519     4%      /var
 /dev/dsk/c0t4d0s5     404363    47419    316514     14%     /opt
 /dev/dsk/c0t4d0s3     384883    106488   239915     31%     /usr/openwin
 swap                  211520    36       211484     1%      /tmp
 /dev/dsk/c4t0d0s6     8365850   9        7529261    1%      /filesys1
 /dev/dsk/c4t0d1s6     4139980   9        3725981    1%      /filesys2
 /dev/dsk/c4t0d2s6     4139980   9        3725981    1%      /filesys3
 /dev/dsk/c4t0d3s6     4139980   9        3725981    1%      /filesys4
 /dev/dsk/c4t0d4s6     4139980   9        3725981    1%      /filesys5
 -hosts                0         0        0          0%      /net
 auto_home             0         0        0          0%      /home
 -xfn                  0         0        0          0%      /xfn
 sc2000:vold(pid232)   0         0        0          0%      /vol
```

df -ak Command (7190 Model 100 to SBus Adapter)

Figure 146. Output of

4. If you have made entries in /etc/vfstab to automatically mount the file systems associated with the SSA drives, comment these entries out. Figure 147 shows our example with the appropriate entries commented out.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
#/dev/dsk/c4t0d0s6  /dev/rdsk/c4t0d0s6  /filesys1  ufs  1  yes  -
#/dev/dsk/c4t0d1s6  /dev/rdsk/c4t0d1s6  /filesys2  ufs  1  yes  -
#/dev/dsk/c4t0d2s6  /dev/rdsk/c4t0d2s6  /filesys3  ufs  1  yes  -
#/dev/dsk/c4t0d3s6  /dev/rdsk/c4t0d3s6  /filesys4  ufs  1  yes  -
#/dev/dsk/c4t0d4s6  /dev/rdsk/c4t0d4s6  /filesys5  ufs  1  yes  -
```

Figure 147. Sample /etc/vfstab File (7190 Model 100 to SBus Adapter)

5. Shut down the system. Install the IBM SSA Interface Controller for Sun SBus in accordance with the *IBM SSA Interface Controller for Sun SBus Installation and User's Guide*.
6. Boot the system and install the SBus driver and utilities package in accordance with the *IBM SSA Interface Controller for Sun SBus Installation and User's Guide*.

7. After the package has been installed and the system rebooted, use the format command to identify the new controller slot number that has been assigned to the SSA drives. In our sample case, the device path names have changed from the 7190 Model 100 nomenclature of c4t*d*s* to the SBus adapter nomenclature of c2t*d0s*. The 7190 Model 100 in the sample case uses LUN priority mapping (mode 2 or 3).

8. Create a new map file for the drives, using the **sscf -d** command. Figure 148 demonstrates this command and its output.

```
# sscf -a 0 -d > mapfile.0
# view mapfile.0
 # IDs of installed Adapters: 0 1
 # Configuration map for Adapter 0
 # UID     UID    SCSI
 # High    Low     ID
   0004  ac7e4fab  00
   0004  ac7e50cb  01
   0004  ac7e4f49  02
   0004  ac7e50a8  03
   0004  ac9de6d8  04
```

sscf to Create a Map File (7190 Model 100 to SBus Adapter)

Figure 148. Using

9. Edit the map file created by the **sscf** command. Assign the desired SCSI target IDs to the respective SSA drives. Figure 149 shows our sample map file edited to reflect the original configuration used on the 7190.

   Note that the 7190 Model 100 maps drives to LUNs 0 through 7 on SCSI target IDs 0 through 5, 6 (in mode 0), and 8 through 15. The SBus controller maps drives to LUN 0 on SCSI target IDs 0 through 47. If this changes your SCSI controller/ID/LUN combinations and your system has any applications that make direct access to device path names, you must update the application references to reflect the updated path names. This task can vary in complexity depending on your particular applications.

```
# IDs of installed Adapters: 0 1
# Configuration map for Adapter 0
# UID     UID    SCSI
# High    Low     ID
  0004  ac9de6d8  00
  0004  ac7e50a8  01
  0004  ac7e4f49  02
  0004  ac7e50cb  03
  0004  ac7e4fab  04
```

Figure 149. Updating SBus Map File to 7190 Model 100 Assignments

10. Use the SSCF utility with the **-p** (permanent) and **-m** (mapfile) options to make your edited map file permanent. For our sample test environment, the command would be:

    sscf -a 0 -p -m mapfile.0

11. Using the output of the **df** command from Step 3 on page 193, along with the updated map file, create a table to show the relationship between the

original 7190 Model 100 device path names and the new SBus adapter path names for the SSA devices. Include the applicable mount points. An example table for our test environment is shown in Figure 150 on page 196.

```
7190 model 100 (Original)    Mount        SBus Interface
Device Names       Points       Device Names
-----------------  -----------  ---------------

/dev/dsk/c4t0d0s6  /filesys1    /dev/dsk/c2t0d0s6
/dev/dsk/c4t0d1s6  /filesys2    /dev/dsk/c2t1d0s6
/dev/dsk/c4t0d2s6  /filesys3    /dev/dsk/c2t2d0s6
/dev/dsk/c4t0d3s6  /filesys4    /dev/dsk/c2t3d0s6
/dev/dsk/c4t0d4s6  /filesys5    /dev/dsk/c2t4d0s6
```

Figure 150. 7190 Model 100 and New SBus Adapter Device Names

12. Using the table you have created, update all applications that use a device path name to reference the respective partitions or mount points. As a minimum, you have to update /etc/vfstab if you are using it to automatically mount partitions. Figure 151 shows an example of the new /etc/vfstab for our sample test environment.

```
#device  device  mount  FS fsck mount mount
#to mount to fsck  point  type pass at boot options
#
fd  -  /dev/fd  fd  -  no  -
/proc  -  /proc  proc  -  no  -
/dev/dsk/c0t4d0s1  -  -  swap  -  no  -
/dev/dsk/c0t4d0s0  /dev/rdsk/c0t4d0s0  /  ufs  1  no  -
/dev/dsk/c0t4d0s6  /dev/rdsk/c0t4d0s6  /usr  ufs  1  no  -
/dev/dsk/c0t4d0s4  /dev/rdsk/c0t4d0s4  /var  ufs  1  no  -
/dev/dsk/c0t4d0s5  /dev/rdsk/c0t4d0s5  /opt  ufs  2  yes  -
/dev/dsk/c0t4d0s3  /dev/rdsk/c0t4d0s3  /usr/openwin  ufs  2  yes  -
swap  -  /tmp  tmpfs  -  yes  -
/dev/dsk/c2t0d0s6  /dev/rdsk/c2t0d0s6  /filesys1  ufs  1  yes  -
/dev/dsk/c2t1d0s6  /dev/rdsk/c2t1d0s6  /filesys2  ufs  1  yes  -
/dev/dsk/c2t2d0s6  /dev/rdsk/c2t2d0s6  /filesys3  ufs  1  yes  -
/dev/dsk/c2t3d0s6  /dev/rdsk/c2t3d0s6  /filesys4  ufs  1  yes  -
/dev/dsk/c2t4d0s6  /dev/rdsk/c2t4d0s6  /filesys5  ufs  1  yes  -
```

Figure 151. Updated /etc/vfstab for 7190 Model 100 to SBus Adapter.

13. Incorporate the updated device names into the kernel configuration. There are two ways to accomplish this. One is to issue this command:

    reboot -- -r

    The other is to issue the following series of commands:

    drvconfig

    devlinks

    disks

    ucblinks

    The method you use depends on your particular system environment.

Appendix A. Special Notices

This publication is intended to help client storage administrators and technical support staff to plan, configure, design, install, and implement SSA disk subsystems attached to Sun Solaris servers that use SCSI host interfaces. The publication shows how to use the IBM SSA Interface Controller for Sun SBus to enable SSA disk subsystem exploitation through an internal adapter card. The publication also demonstrates how to use the 7190 SCSI Host to SSA Loop Attachment product to implement SSA disk subsystems using an external converter product. The information in this publication is not intended as the specification of any programming interfaces that are provided by the IBM SSA Interface Controller for Sun SBus and the 7190 SCSI Host to SSA Loop Attachment. See the PUBLICATIONS section of the IBM Programming Announcement for the IBM SSA Interface Controller for Sun SBus and the 7190 SCSI Host to SSA Loop Attachment 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

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 pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

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.

The following document contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples contain the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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

AS/400
IBM
RS/6000
Ultrastar
BookManager
PROFS
System/390

The following terms are trademarks of other companies:

C-bus is a trademark of Corollary, Inc.

SunSBus, Solaris, Java,
and HotJava are trademarks of Sun Microsystems, Incorporated.

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

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

Pentium, MMX, ProShare, LANDesk, and ActionMedia are trademarks or registered trademarks of Intel Corporation in the U.S. and other countries.

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

Other company, product, and service names may be trademarks or service marks of others.

Appendix B. Related Publications

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

B.1 International Technical Support Organization Publications

For information on ordering these ITSO publications see “How to Get ITSO Redbooks” on page 201.

- *A Practical Guide to Serial Storage Architecture for AIX* , SG24-4599
- *AIX Storage Management*, GG24-4484
- *Monitoring and Managing IBM SSA Disk Subsystems*, SG24-5251 (available at a later date)

B.2 Redbooks on CD-ROMs

Redbooks are also available on CD-ROMs. **Order a subscription** and receive updates 2-4 times a year at significant savings.

CD-ROM Title	Subscription Number	Collection Kit Number
System/390 Redbooks Collection	SBOF-7201	SK2T-2177
Networking and Systems Management Redbooks Collection	SBOF-7370	SK2T-6022
Transaction Processing and Data Management Redbook	SBOF-7240	SK2T-8038
Lotus Redbooks Collection	SBOF-6899	SK2T-8039
Tivoli Redbooks Collection	SBOF-6898	SK2T-8044
AS/400 Redbooks Collection	SBOF-7270	SK2T-2849
RS/6000 Redbooks Collection (HTML, BkMgr)	SBOF-7230	SK2T-8040
RS/6000 Redbooks Collection (PostScript)	SBOF-7205	SK2T-8041
RS/6000 Redbooks Collection (PDF Format)	SBOF-8700	SK2T-8043
Application Development Redbooks Collection	SBOF-7290	SK2T-8037

How to Get ITSO Redbooks

This section explains how both customers and IBM employees can find out about ITSO redbooks, CD-ROMs, workshops, and residencies. A form for ordering books and CD-ROMs is also provided.

This information was current at the time of publication, but is continually subject to change. The latest information may be found at http://www.redbooks.ibm.com/.

How IBM Employees Can Get ITSO Redbooks

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

- **Redbooks Web Site on the World Wide Web**

  http://w3.itso.ibm.com/

- **PUBORDER** — to order hardcopies in the United States
- **Tools Disks**

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

  ```
  TOOLCAT REDPRINT
  TOOLS SENDTO EHONE4 TOOLS2 REDPRINT GET SG24xxxx PACKAGE
  TOOLS SENDTO CANVM2 TOOLS REDPRINT GET SG24xxxx PACKAGE (Canadian users only)
  ```

  To get BookManager BOOKs of redbooks, type the following command:

  ```
  TOOLCAT REDBOOKS
  ```

  To get lists of redbooks, type the following command:

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

  To register for information on workshops, residencies, and redbooks, type the following command:

  ```
  TOOLS SENDTO WTSCPOK TOOLS ZDISK GET ITSOREGI 1998
  ```

- **REDBOOKS Category on INEWS**
- **Online** — send orders to: USIB6FPL at IBMMAIL or DKIBMBSH at IBMMAIL

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (http://www.redbooks.ibm.com/redpieces.html). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

How Customers Can Get ITSO Redbooks

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

- **Online Orders** — send orders to:

	IBMMAIL	**Internet**
In United States:	usib6fpl at ibmmail	usib6fpl@ibmmail.com
In Canada:	caibmbkz at ibmmail	lmannix@vnet.ibm.com
Outside North America:	dkibmbsh at ibmmail	bookshop@dk.ibm.com

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	(long distance charges apply)
(+45) 4810-1320 - Danish	(+45) 4810-1020 - German
(+45) 4810-1420 - Dutch	(+45) 4810-1620 - Italian
(+45) 4810-1540 - English	(+45) 4810-1270 - Norwegian
(+45) 4810-1670 - Finnish	(+45) 4810-1120 - Spanish
(+45) 4810-1220 - French	(+45) 4810-1170 - Swedish

- **Mail Orders** — send orders to:

IBM Publications Publications Customer Support P.O. Box 29570 Raleigh, NC 27626-0570 USA	IBM Publications 144-4th Avenue, S.W. Calgary, Alberta T2P 3N5 Canada	IBM Direct Services Sortemosevej 21 DK-3450 Allerød Denmark

- **Fax** — send orders to:

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	(+45) 48 14 2207 (long distance charge)

- **1-800-IBM-4FAX (United States)** or **(+1)001-408-256-5422 (Outside USA)** — ask for:

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

- **On the World Wide Web**

Redbooks Web Site	http://www.redbooks.ibm.com/
IBM Direct Publications Catalog	http://www.elink.ibmlink.ibm.com/pbl/pbl

Redpieces

For information so current it is still in the process of being written, look at "Redpieces" on the Redbooks Web Site (`http://www.redbooks.ibm.com/redpieces.html`). Redpieces are redbooks in progress; not all redbooks become redpieces, and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbook Order Form

Please send me the following:

Title	Order Number	Quantity

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

- Invoice to customer number
- Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

Numerics

A

B

C

D

E

F

I

L

M

O

P

R

S

T

U

V

ITSO Redbook Evaluation

Exploiting SSA Disk Subsystems in Sun Solaris Platform Environments
SG24-5083-00

Your feedback is very important to help us maintain the quality of ITSO redbooks. **Please complete this questionnaire and return it using one of the following methods:**

- Use the online evaluation form found at http://www.redbooks.ibm.com
- Fax this form to: USA International Access Code + 1 914 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Which of the following best describes you?
__Customer __Business Partner __Solution Developer __IBM employee
__None of the above

Please rate your overall satisfaction with this book using the scale:
(1 = very good, 2 = good, 3 = average, 4 = poor, 5 = very poor)

Overall Satisfaction _____

Please answer the following questions:

Was this redbook published in time for your needs? Yes____ No____

If no, please explain:

_____

_____

_____

_____

What other redbooks would you like to see published?

_____

_____

_____

Comments/Suggestions: **(THANK YOU FOR YOUR FEEDBACK!)**

_____

_____

_____

_____

_____

SG24-5083-00
Printed in the U.S.A.