

Sun StorEdge™ A7000 DataShare Open Backup System Administrator's Guide



THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No. 805-6520-10
December 1998, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 USA. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook, Java, the Java Coffee Cup, StorEdge, Sun Enterprise NetBackup, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook, Java, le logo Java Coffee Cup, StorEdge, Sun Enterprise NetBackup, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

Preface	xvii
Before You Read This Book	xvii
How This Book Is Organized	xviii
Using UNIX Commands	xviii
Documentation Conventions	xix
Shell Prompts	xx
Related Documentation	xx
Sun Documentation on the Web	xx
Sun Welcomes Your Comments	xxi
1. Product Description and A7000 Architecture	1-1
DataShare Open Backup	1-2
MVS Mainframes as the Backup Server	1-3
Sun as the Backup Server	1-4
StorEdge A7000 Architecture	1-6
DataShare Architecture	1-7
What DataShare Open Backup Provides	1-8
What DataShare Open Backup Does Not Provide	1-8

2. Using the Mainframe as the Backup Server	2-1
Introduction	2-2
Terminology	2-3
Example System	2-3
Assumptions	2-4
Planning Process	2-5
Determining the Size and Number of Open Disk Volumes to be Backed Up	2-5
Selecting the Type and Number of Mainframe Volumes to be Used as Open System Host Volumes	2-6
Deciding How to Connect the Open Systems Host Volumes to the Mainframe	2-8
Choosing a Naming Convention for the DataShare Open Backup Host Volumes and DataShare Open Backup Volume Datasets	2-9
Data Protection Considerations	2-10
Determining the Required Esoteric Device Name Changes or Additions	2-11
Scheduling a Time with Sun Enterprise Services for Implementation of Configuration Changes	2-11
Implementation	2-13
Updating the MVS I/O Configuration	2-13
Varying the Mainframe Volumes Offline	2-13
Updating the DataShare Open Backup Configuration	2-13
Restarting the System	2-14
Initializing	2-14
Using DataShare Open Backup in Backup and Restore Jobs	2-14
Quiescing the Client System	2-15
Optimizing Backup and Restore Performance	2-15
JCL Examples	2-16
Example 1	2-16

Example 2 2-18

Example 3 2-20

Automatically Coordinating Processing With the Open System 2-21

Implications for Business Continuation Planning 2-22

Planning Worksheet 2-23

3. Using the Sun System as the Backup Server 3-1

Introduction 3-2

Terminology 3-2

Planning 3-4

Capacity 3-4

Example – Determining the Bandwidth Necessary to Back Up Data
Within the Backup Window 3-4

Frequency 3-5

Performance 3-5

Recovery 3-5

Setup 3-7

▼ To Set Up an A7000 Client Volume for Use With the Sun Backup Server 3-9

Using DataShare Open Backup 3-10

▼ To Back Up Data 3-10

▼ To Restore Data 3-11

Planning Worksheet 3-12

4. Client Issues 4-1

Introduction 4-2

Disk Partitions 4-3

Device Types 4-6

Changing Disk Hardware Configurations 4-6

Restoring Elsewhere 4-6

Sun Solaris Operating Environment 2.5.1 and 2.6 4-7

Disk Naming Conventions 4-7

- ▼ To Mount a File System 4-8
- ▼ To Unmount a File System 4-8
- ▼ To Backup a Sun System 4-9
- ▼ To Restore a Sun System 4-10

Windows NT 4.0 4-11

Disk Naming Conventions 4-11

- ▼ To Mount a File System 4-11
- ▼ To Unmount a File System 4-11
- ▼ To Backup a Windows NT System 4-12
- ▼ To Restore a Windows NT System 4-13
 - ▼ To Write a Signature to the New Disk 4-14
 - ▼ To Label the New Disk 4-14

HP-UX 10.01 and 10.20 4-15

Disk Naming Conventions 4-15

- ▼ To Mount a File System 4-16
- ▼ To Unmount a File System 4-16
- ▼ To Backup an HP-UX System 4-16
- ▼ To Restore an HP-UX System 4-17

IBM AIX 4.1.2, 4.1.3, and 4.2 4-19

Disk Naming Conventions 4-19

- ▼ To Mount a File System 4-19
- ▼ To Unmount a File System 4-20
- ▼ To Backup an IBM AIX System 4-20
- ▼ To Restore an IBM AIX System 4-21

5. Best Business Practices	5-1
Introduction	5-2
Mainframe Backup of Open Systems Data	5-2
Current Method	5-2
Mainframe Environment	5-2
Open Systems Environment	5-3
Application Details	5-3
Mainframe Hardware and Configuration	5-3
Open Systems Hardware and Configuration	5-3
Problem Analysis	5-4
Solution Description	5-4
Proposed Method of Accomplishing Solution	5-4
▼ To Do Capacity Planning	5-4
▼ To Configure the A7000 (Completed by Sun Enterprise Services Personnel)	5-5
▼ To Configure the Mainframe	5-5
▼ To Configure the Open Systems	5-6
Application Solution	5-6
Backup Considerations	5-6
Restore Considerations	5-6
Sun Backup of Open Systems Data	5-7
Current Method	5-7
Application Details	5-7
Sun Hardware and Operating System	5-7
Open Systems Hardware and Operating System	5-8
Problem Analysis	5-8
Solution Description	5-8
Capacity Planning	5-8

Configuring the Open Systems 5-9

Application Solution 5-9

Backup Considerations 5-9

Restore Considerations 5-9

A. A7000 Platform-Specific Storage Information A-1

Sun Solaris Operating Environment 2.5.1 and 2.6 A-2

▼ To Scan or Rescan the Hardware A-2

▼ To Format and Partition Disks A-2

▼ To Verify the Disk Partitioning A-3

▼ To Make a UNIX File System (UFS) A-3

Mounting File Systems and Starting Volumes A-3

▼ To Use the Volume Manager (VxVM) A-4

▼ To Use a UNIX File System A-4

Unmounting File Systems and Stopping Volumes A-5

▼ To Stop the Volume Manager (VxVM) A-5

▼ To Unmount a UNIX File System A-5

Windows NT 4.0 A-6

▼ To Obtain SCSI IDs A-6

Displaying the Configured Disks A-7

▼ To Login in as the Disk Administrator A-7

Creating a Windows NT File System A-8

▼ To Create a Primary Partition in the Disk Administrator A-8

▼ To Format the Partition A-8

HP-UX 10.01 and 10.20 A-10

▼ To Scan or Rescan the Hardware A-10

Creating an HP-UX File System A-11

▼ To Format and Partition the Disk and Create a File System Using the System Administration Manager (SAM) A-11

▼ To Format and Partition the Disk and Create a File System From the Command Line (Without SAM) A-13

IBM AIX 4.1.2, 4.1.3, and 4.2 A-18

▼ To Scan or Rescan the Hardware A-18

Creating the IBM AIX File System A-19

▼ To Optimize Device Attributes A-19

▼ To Use the SMIT to Configure Sun Disks A-20

B. A7000 Configurations and Limitations B-1

Example 1: Maximum A7000 Configuration B-2

Available External I/O Connection Types B-2

324 Disk Drives (Internal) B-2

14 Slots Per Node B-2

Example 2: A7000 Configuration B-3

32 External I/O Connections (16 per Node) B-3

108 Internal Disk Drives B-3

14 Slots Per Node B-3

Example 3: A7000 Configuration B-4

32 external I/O connections (16 per node) B-4

240 Disk Drives (Internal) B-4

14 Slots per node B-4

Glossary Glossary-1

Index Index-1

Code Samples

CODE EXAMPLE 2-1 Backup with FDR 2-17

CODE EXAMPLE 2-2 Restore with FDR 2-18

CODE EXAMPLE 2-3 Backup with DFSMSdss 2-19

CODE EXAMPLE 2-4 Restore with DFSMSdss 2-19

CODE EXAMPLE 2-5 Backup with IEBGENER 2-20

CODE EXAMPLE 2-6 Restore with IEBGENER 2-21

CODE EXAMPLE 2-7 Automatically Coordinating Processing With the Open System 2-21

Figures

- FIGURE 1-1 DataShare Open Backup 1-3
- FIGURE 1-2 MVS Mainframe as the Backup Server 1-4
- FIGURE 1-3 Sun Solaris as the Backup Server 1-6
- FIGURE 1-4 Product Architecture 1-7
- FIGURE 1-5 DataShare Architecture 1-8
- FIGURE 2-1 Mainframe as Backup Server 2-2
- FIGURE 2-2 DataShare Open Backup Large Volume Mapping 2-6
- FIGURE 2-3 DataShare Open Backup Smaller Volume Mapping 2-7
- FIGURE 3-1 DataShare Open Backup with a Sun Backup Server 3-3
- FIGURE 3-2 Phantom Header and Tail Usage 3-7
- FIGURE 3-3 Two Open System Clients and the Sun Server Connected to the A7000 3-8
- FIGURE 4-1 Relationship Between A7000, Backup Server, and Open System Client 4-2
- FIGURE 4-2 Single Volume 4-3
- FIGURE 4-3 Three Volumes on a Disk 4-4
- FIGURE 4-4 Two Disks in a Single Volume 4-4
- FIGURE 4-5 Data as a Single Disk 4-5
- FIGURE 4-6 LUN Composed of a Virtual Partition 4-5

Tables

TABLE P-1	Documentation Conventions	xix
TABLE P-2	Shell Prompts	xx
TABLE P-3	Related Documentation	xx
TABLE 2-1	UNIX Example System Volumes	2-3
TABLE 2-2	Data Capacity	2-7

Preface

The *DataShare Open Backup System Administrator's Guide* describes the Sun StorEdge™ A7000 DataShare Open Backup product and provides planning, setup, and usage information.

This document is designed for experienced administrators of mainframes and Sun systems who will plan and set up their systems to do backups and restores using DataShare Open Backup with their existing mainframe and Sun backup and restore software, peripherals, and procedures directly to A7000 resident open system volumes.

The audience for this manual consists of the following groups:

- Mainframe administrators (system programmers)
- Open system administrators

Before You Read This Book

You must have thorough knowledge of your mainframe and open systems platforms to fully use the information in this document. Refer to your vendor documentation.

How This Book Is Organized

Chapter 1 describes the DataShare Open Backup product.

Chapter 2 contains planning, setup, and usage information when using an MVS mainframe as a server.

Chapter 3 contains planning, setup, and usage information when using a Sun system as a server.

Chapter 4 describes client platform issues.

Chapter 5 contains best business practices for specific configurations.

Appendix A contains basic information and procedures that enable you to set up your supported open system platforms for use with the A7000.

Appendix B contains Sun StorEdge A7000 configuration and limitation information.

Glossary includes definitions of DataShare Open Backup, mainframe, and open system terms.

Using UNIX Commands

This document may not contain information on basic UNIX[®] commands and procedures such as shutting down the system, booting the system, and configuring devices.

Refer to the UNIX documentation that you received with your system

Documentation Conventions

TABLE P-1 Documentation Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output.	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be <code>root</code> to do this. To delete a file, type <code>rm filename</code> .
[]	In syntax, brackets indicate that an argument is optional.	In this example, <code>/etc/dsfadmin</code> is the only part of the syntax that <i>must</i> be entered. <code>/etc/dsfadmin [-v [device]]</code>
com(n)	The form <code>command(number)</code> , where the number in parentheses ranges from 1 through 8 and is followed by letters, indicates the presence of an online reference man page.	<code>dsfadmin(1DSF)</code>

Shell Prompts

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

TABLE P-3 Related Documentation

Application	Title	Part Number
Release	<i>DataShare Open Backup Release Notes</i>	online
Reference	<i>Direct Access Storage Device (DASD) Manager User's Guide</i>	805-4884-10
Reference	<i>SCSI Target Emulation Release Notes</i>	online
Reference	<i>Simulation of Count-Key-Data (SIMCKD) Release Notes</i>	online

Sun Documentation on the Web

The `docs.sun.comsm` web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

`http://docs.sun.com`

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

`docfeedback@sun.com`

Please include the part number of your document in the subject line of your email.

Product Description and A7000 Architecture

This chapter describes DataShare Open Backup and how it relates to the StorEdge A7000 architecture.

- DataShare Open Backup—page 1-2
- StorEdge A7000 Architecture—page 1-6
- DataShare Architecture—page 1-7
- What DataShare Open Backup Provides—page 1-8
- What DataShare Open Backup Does Not Provide—page 1-8

DataShare Open Backup

DataShare Open Backup allows you to use standard software packages executing on Sun systems or MVS or VSE mainframe hosts for full volume level backups and restores of UNIX and Windows NT data that is on the StorEdge A7000. Consolidating the backup of your open system clients to one backup server frees host cycles and leverages software and peripherals across multiple platforms.

DataShare Open Backup presents the volumes targeted for backup and restore operations in the format expected by the Backup Host (server). This allows the backup server to use its standard backup utilities to perform these volume level operations for attached hosts (clients).

Note – Only full volume backup and restore are supported at this time.
The supported backup utility on VSE is FCOPY.




FIGURE 1-1 DataShare Open Backup

MVS Mainframes as the Backup Server

When an MVS mainframe is the backup server for DataShare Open Backup, the open system resident volumes appear as datasets in one or more 3380-K or 3390-3 volumes.

The mainframe backup server can be any IBM or plug-compatible mainframe running MVS/ESA 4.x or compatible versions, or OS/390. You can use any MVS mainframe backup and restore product that supports 3380-K or 3390-3 datasets (for example, Fast Dump Restore (FDR) and Data Facility Data Set Services (DFDSS)) to back up and restore the A7000-resident open system full volumes.




FIGURE 1-2 MVS Mainframe as the Backup Server

Additional software enhancing the automation of the backup and recovery process can include automated job scheduling and tape library management software.

Back end storage can consist of:

- 3490 compatible tape drives (IBM 3590 Magstar)
- Automated tape libraries
- Virtual tape systems

The A7000 provides 3990-3 control unit function and up to 512 logical volume images in 3380-A/D/J/K or 3390-1/2/3/9 format. DataShare Open Backup presentations of open volumes are limited to 3380-K or 3990-3.

You can use a maximum of four ESCON or parallel channels per MVS or OS/390 image to connect the mainframe to the A7000. Up to 32 host interface connections are available to the A7000, which can be divided in units of four between SCSI (STE) and mainframe ESCON or parallel channels.

Sun as the Backup Server

When a Sun system is the backup server for DataShare Open Backup, the open system resident volumes appear as raw volumes to the backup server. You can use any Sun Solaris operating environment utility capable of backing up and restoring raw volume images (for example, Sun™ Enterprise NetBackup™ or `dd`) to back up and restore the A7000-resident open system full volumes.

Note – DataShare Open Backup does not support Sun backup of MVS or VSE mainframe volumes.




FIGURE 1-3 Sun Solaris as the Backup Server

StorEdge A7000 Architecture

The StorEdge A7000 architecture includes the Simulation of Count-Key-Data (SIMCKD) and SCSI Target Emulation (STE) facilities, which allow the simultaneous emulation of Count-Key-Data (CKD) mainframe volumes using Block Mux Channel (BMC) and Enterprise Systems Connection (ESCON) Channels, and Fixed Block Architecture (FBA) open systems volumes across SCSI II Channels. SIMCKD and STE share a common cache and a head/disk assembly (HDA) pool within the A7000. For a specified volume, information is physically stored on the HDAs and staged into cache in a format that is optimized for the native emulation being provided.

The following figure illustrates the A7000 architecture and the relationship between the cache, HDA pool, and SIMCKD and STE emulation facilities.




FIGURE 1-4 Product Architecture

DataShare Architecture

DataShare is an integrated function within the A7000 architecture that is configured to access the HDA pool providing non-native information through the cache to an emulation different from the one that created it. The STE layer is used to make data available to open system hosts and the SIMCKD layer is used to make data available to mainframe hosts.

DataShare acts directly on stored data in real time as a non-native host reads information. Host platform configuration of non-native devices needing to be shared is consistent with the native emulation (for example, STE for open systems). File system appearance for non-native information is consistent with a native supported file system for the host requiring access. The following figure illustrates the DataShare architecture.




FIGURE 1-5 DataShare Architecture

What DataShare Open Backup Provides

DataShare Open Backup is an effective answer for many existing data center problems and also provides new opportunities for operational efficiency. DataShare Open Backup provides these features:

- Enables a business to recover data in the event of a disaster, allowing continuation of normal operation.
- Consolidates client system I/O peripherals.
- Offloads machine cycles (backup process overhead) from the clients to the shared backup server.
- Uses the advanced features of the host platform to provide backup and recovery for the open system client platform (for example, mainframe tape management).
- Minimizes the need to invest in new open systems backup and restore capabilities.
- Applies data center class backup and recovery technology to the entire enterprise.
- Eliminates the need for new procedures and training for each client platform.
- Restores a system that is not operational.
- Year 2000 compliance. This means that DataShare Open Backup (a) will not produce errors in date data related to the year change from December 31, 1999 to January 1, 2000 and (b) will handle leap years correctly.

The two digit year field in a mainframe file date is interpreted as follows before the year information is placed in the ISO-9660 image:

- Based on the year 2000 if the number is in the range of 00 to 68
- Based on the year 1900 if the number is in the range of 69 to 99

What DataShare Open Backup Does Not Provide

DataShare Open Backup does not provide these functions:

- File level backup
- Backup of volumes while mounted for write-access
- Actual backup and restore software to be run on the backup server
- Sun backup of mainframe volumes

Using the Mainframe as the Backup Server

This chapter describes how to plan and implement a mainframe running MVS/ESA or OS/390 as a backup server for a DataShare Open Backup configuration.

- Introduction—page 2-2
- Assumptions—page 2-4
- Planning Process—page 2-5
- Implementation—page 2-13
- Using DataShare Open Backup in Backup and Restore Jobs—page 2-14
- Planning Worksheet—page 2-23

Introduction

DataShare Open Backup makes open systems volumes on an A7000 available to your existing mainframe backup and restore tools by presenting the open system volumes to the mainframe as fixed block QSAM datasets allocated within standard 3380-K or 3390-3 volumes. The following figure illustrates this concept.




FIGURE 2-1 Mainframe as Backup Server

After DataShare Open Backup is configured, use it by including A7000-resident volumes or datasets in your normal backup and recovery procedures.

Terminology

The following terms are used throughout this chapter. Refer to the Glossary for definitions of mainframe and open systems terms.

Term	Definition
Backup window	Time available to back up a quantity of data based on the availability requirements of the business
Client	Open System Client machine. That is, an HP, NT, AIX or Sun
Esoteric	Identifier by which a group of devices (usually tape and disk) are known
Open system host volume	3380 or 3390 volume containing the open system volume image
Open system host volume dataset	QSAM dataset on the open system host volume in which that image resides

Example System

TABLE 2-1 contains the volumes requiring backup on a UNIX system that is used as an example throughout this chapter to illustrate various concepts and syntax.

TABLE 2-1 UNIX Example System Volumes

Volumes Requiring Backup
40.0-Gbyte database (raw) in five 9-Gbyte volumes
4.8-Gbyte other data (file system) in one 4.8-Gbyte volume
1.0-Gbyte boot volume (root file system) in one 1-Gbyte volume
45.8-Gbyte Total

Assumptions

Certain assumptions are made in this chapter:

- You are currently using at least one of the backup and restore products available for MVS/ESA or OS/390. Any utility that can minimally copy a disk dataset to tape and back again can be used, including standard DFSMSdfp utility programs.
- You are familiar with the Hardware Configuration Definition (HCD) process on MVS, you know how to add a 3990 compatible storage director to your system, and you can set up and maintain strings of standard 3380 and 3390 compatible storage devices.
- Because DataShare Open Backup does only full volume backups, your data does not require incremental backup. If the data does require incremental backups, you are supplementing DataShare Open Backup with a local incremental backup process.
- The time required to back up and restore the open system host volumes is similar to that required to back up and restore any other equivalent mainframe volume. These times can vary according to the concurrent system workload, the speed and capacity of the tape subsystem, the number and type of channel connections to both the A7000 and the tape subsystem, and the concurrent I/O workload on these channels.

Planning Process

Successful planning for DataShare Open Backup requires:

- Determining the size and number of open disk volumes to be backed up
- Selecting the type and number of mainframe volumes to be used to represent open system host volumes
- Determining the best way to connect the DataShare Open Backup volumes to the mainframe
- Choosing a naming convention for the DataShare Open Backup host volumes and DataShare Open Backup volume datasets
- Determining the required esoteric device name changes or additions
- Scheduling a time with Sun Enterprise Services to implement the configuration changes

Use this information to modify your mainframe I/O configuration (if necessary). Sun Enterprise Services personnel use this information to configure the A7000. Use the worksheet at the end of this chapter to help with your planning.

Determining the Size and Number of Open Disk Volumes to be Backed Up

The term *open disk volume* normally equates to a SCSI target logical unit number (LUN) from the open system host viewpoint. This concept may not be clear in complex environments. See Chapter 4 for more detailed open systems client-specific information. Consult Sun Enterprise Services if there is any ambiguity.

Because open systems volumes must be quiescent during the backup, determine the timing and frequency of backups by the availability and duration of off-peak client usage periods. Choosing nonbusiness or off-peak usage hours lessens the chance of adversely affecting client I/O response time.

The database in the example system (TABLE 2-1) must be quiescent for the duration of the backup. Design your configuration so that backup is complete before the database must be restarted.

The example boot volume is relatively small, but the entire client system must be shut down to use DataShare Open Backup to do the backup.

In the example in this chapter, different volumes on the open system client can require different backup strategies. For example, the database must be backed up nightly and the boot volume only once a month. These considerations may affect MVS I/O configuration choices.

Selecting the Type and Number of Mainframe Volumes to be Used as Open System Host Volumes

Each open disk volume must be mapped individually to one or more datasets on either 3380-K or 3390-3 mainframe volumes.

Choose 3390-3 unless you are constrained by your existing I/O configuration. For example, you need to use a range of addresses already reserved, or the volumes are being added to an A7000 already being used for mainframe storage.

The open volume space must be allocated as several discrete datasets. One volume can contain more than one dataset with a maximum of eight datasets per volume, but datasets cannot span volumes. Open system volumes that are larger than a mainframe volume are divided into several individual datasets on separate volumes. The following figure illustrates large volume mapping.




FIGURE 2-2 DataShare Open Backup Large Volume Mapping

Smaller open system volumes can be represented as different datasets on the same Direct Access Storage Device (DASD) volume if the DASD device has room. The following figure illustrates smaller volume mapping.




FIGURE 2-3 DataShare Open Backup Smaller Volume Mapping

As a rule of thumb for capacity calculations, a 3380-K can store approximately 1.7 Gbytes of data and a 3390-3 can store approximately 2.5 Gbytes. TABLE 2-2 shows the available data capacity using the actual blocking factors of DataShare Open Backup.

TABLE 2-2 Data Capacity

TYPE	Cylinders	Tracks/ CYL	Block Size	Blocks per track	Data Capacity per Volume
3390-3	3338	15	27K	2	2.578 Gbytes
3380-K	2654	15	22.5K	2	1.708 Gbytes

Leftover space on the mainframe volume is not wasted physical space on the A7000. The 3380-K or 3390-3 is a logical *view* of the data on an underlying *real* open system volume. Space on the real volume is not wasted.

Map the open system data using as few DASD addresses as possible. The maximum addressable capacity for any A7000 from a mainframe is based on the largest possible number of device addresses and the capacity of each device. If you use 3380-K devices, or if you use devices that contain a lot of leftover space, you may exhaust the capacity of the mainframe to address the data before you exhaust the physical capacity of the A7000.

To illustrate capacity calculations, see TABLE 2-1. In that example, a total of 45.8 Gbytes of space must be allocated. Most of the data (40 Gbytes) is on 9-Gbyte volumes. Each 9-Gbyte volume requires four volumes each when using 3390-3s, or six volumes each when using 3380-Ks.

Using 3390-3s, the example system requires:

- 20 mainframe volumes for the 9-Gbyte volumes
- Two mainframe volumes for the 4.8-Gbyte volume
- One mainframe volume for the boot volume
- Total of 23 mainframe volumes

If you use 3380-Ks, you must have 35 volumes. If your configuration has only 32 available addresses, you must use 3390s even though either allocation uses the same physical space on the A7000.

Consider whether to reserve any spare volumes to use for individual file recovery. While it is not possible to do a file level backup using DataShare Open Backup, you can recover individual files. Do this by restoring a file system volume to a spare volume, mounting it on a different mount point, and then copying the files from the restored file system to the desired one.

Only backups created from volumes containing standard file systems can be handled in this way. In the chapter example, reserve a 4.8-Gbyte volume as a spare. This volume is suitable to restore either the 4.8-Gbyte volume or the boot volume for file level recovery. Do not allocate any space for the database, as it does not contain a file system.

Deciding How to Connect the Open Systems Host Volumes to the Mainframe

DataShare Open Backup can be set up on a variety of A7000 configurations. For example:

- Only open systems storage
- Mainframe and open systems storage
- Expanded from mainframe-only storage to accommodate both types of storage
- Expanded from open system only storage to accommodate both types of storage

The A7000 supports a maximum of 32 host interface connections that can be distributed among mainframe channels (parallel or ESCON) or SCSI channels. How you connect the mainframe to the A7000 affects how you can connect it to the open systems and vice versa. Balance the physical resource and capacity needs of the mainframe and open systems hosts, as well as the service level goals of the system workloads.

See Appendix B for information on how to physically allocate the A7000 resources.

After you have decided the number and type of channels available for the mainframe connection, setting up the 3380 or 3390 volumes to be used for DataShare Open Backup is the same as setting up any other 3990 compatible storage controller and one or more strings of DASD. The DataShare Open Backup volumes require no special handling, and look like other volumes of the same type.

If possible, use additional channels or ESCON rather than parallel channels to obtain better performance. If the backup window is tight, consider backing up the data concurrently from more than one system.

The open system host volumes have special requirements in a System Managed Storage (SMS) environment. New datasets cannot be allocated on them, nor can the open systems volume datasets be migrated to other volumes or media. When processing the data for dump or backup selection, the open system host must be quiescent. Exclude these volumes from any automated SMS handling.

Choosing a Naming Convention for the DataShare Open Backup Host Volumes and DataShare Open Backup Volume Datasets

Sun service personnel assign the names you choose to the DataShare Open Backup host volumes and the DataShare Open Backup volume datasets when the A7000 is configured for DataShare Open Backup.

Only service personnel can change these names. Choose the names carefully, as it is inconvenient to change them later.

If possible, choose volume names that can be related to the underlying open system entity.

The names should be meaningful, relate clearly to the open systems entities they represent, and fit within the limits of your security environment. See the next section, “Data Protection Considerations” for security information.

Using the chapter example system, you might choose the following names:

Volume	Dataset name	Description
BKUP00	OPENBU.DBVOL01.SEQ1	DB volume 1, 1st 3390
BKUP01	OPENBU.DBVOL01.SEQ2	DB volume 1, 2nd 3390
BKUP02	OPENBU.DBVOL01.SEQ3	DB volume 1, 3rd 3390
BKUP03	OPENBU.DBVOL01.SEQ4	DB volume 1, 4th 3390
BKUP04	OPENBU.DBVOL02.SEQ1	DB volume 2, 1st 3390
.		
.		
.		
BKUP19	OPENBU.DBVOL05.SEQ4	DB volume 5, 4th 3390
BKUP20	OPENBU.APPVOL.SEQ1	Other volume, 1st 3390
BKUP21	OPENBU.APPVOL.SEQ2	Other volume, 2ND 3390
BKUP22	OPENBU.ROOTVOL	Boot volume, last 3390
BKUP22	OPENBU.PAGEVOL	Page volume, last 3390

Though it is not required, if the high level qualifier you chose is a user catalog alias, use ISPF 3.4 or IDCAMS DEFINE NONVSAM to catalog these datasets after the devices are online.

A half track blocking factor to the nearest 512 bytes is used for the DataShare Open Backup volume datasets. (The 512-byte increment is required by the underlying disk geometry.) This affords maximum capacity for each device type while providing good performance. These blocksizes are 23040 (22.5K) for a 3380 and 27648 (27K) for a 3390 and are the default.

In summary, the datasets have the following JCL characteristics:

```
DSN=unique.dataset.name
DISK=(OLD,KEEP)
UNIT=3380|3390
VOL=SER=volname
RECFM=F,DSORG=PS,BLKSIZE=23040|27648
```

Data Protection Considerations

As a safeguard, the mainframe cannot modify the Volume Table of Contents (VTOC) of DataShare Open Backup host volumes. Thus, the datasets cannot be accidentally deleted, extended, or reallocated. Similarly, any spare space that appears to be present on the open systems host volumes cannot be used by mainframe applications.

Mount the devices with a private use attribute in your `SYS1.PARMLIB(VATLSTxx)` member, either explicitly or by default. If the devices are mounted public or storage, jobs can fail when they try to allocate files on these volumes.

Although allocation of the DataShare Open Backup volume datasets cannot be modified, the datasets can be overwritten because they must be restorable. If they are overwritten by something other than a restore job, or if the open system is not quiescent, data can be lost or the attached open system can crash.

Note – You can use RACF or an equivalent external security manager to protect the open system volume datasets from unauthorized modification.

Determining the Required Esoteric Device Name Changes or Additions

Because DataShare Open Backup volumes cannot be used for any other purpose, you may want to exclude them from the common esoteric groups such as SYSDA, DISK, WORK, and so forth.

Refer to the previous section, “Data Protection Considerations” for security information.

You can also assign an esoteric name to DataShare Open Backup volumes (for example, `OPNSYSBU`). Though this is not necessary, as the devices can be referenced by their unit number in JCL or by the IBM supplied names 3380, 3390, or `SYSALLDA`, it may be convenient to aggregate them into a single name.

Scheduling a Time with Sun Enterprise Services for Implementation of Configuration Changes

A period of downtime may be required to implement configuration changes. The duration of the downtime depends on the extent of the hardware changes being implemented at a given time. Many installations consolidate changes to minimize the number of outages.

Service personnel change the A7000 software to implement the mainframe storage director and DASD function. The software must be quiescent to allow the changes to be made. Severe error conditions occur on MVS if the volumes are not offline while this is done. At a minimum, all mainframe volumes connected to the A7000 must be varied offline while the service personnel enter and activate the configuration. If this

A7000 is used only as an open systems storage controller, mainframe downtime may not be required. Conversely, if this is an existing mainframe A7000, or if significant hardware changes are needed, such as channel reconfiguration or recabling, an initial program load (IPL) or power-on reset (POR) may also be required.

Note – Any workloads that allocate volumes on the A7000 must be quiescent to enter and implement the A7000 configuration changes.

For example, CICS files residing on volumes on this A7000 must be closed or the CICS region shut down.

In some cases, the system may need an IPL if critical volumes are located on this A7000. Such volumes might include:

- System resident volume
- Volumes containing MVS catalogs
- Volume containing the JES2 checkpoint dataset

When the A7000 is to be shared by both mainframe and open systems, making careful choices of the mainframe data to be kept on it can alleviate this problem. For example, some good choices would include work volumes, one or two local page volumes, or distribution libraries (DLIB). Any volumes that are seldom used or used only for part of the day are also good candidates.

This applies only when configuration changes are first done on the A7000 by service personnel. If the new configuration is considered a long-term, stable one, no long term restrictions on mainframe usage of available volumes are necessary.

Implementation

DataShare Open Backup is implemented by you and Sun service personnel during a scheduled service call.

Note – An IPL or POR may be necessary depending on how extensively the I/O configuration is being altered and whether new equipment is being installed at this time.

Updating the MVS I/O Configuration

- Use the data gathered in the planning process to update your MVS I/O configuration.
- Activate the I/O configuration if it does not conflict with your existing configuration, or have the I/O configuration ready if it cannot be activated before implementing the new configuration.
- You cannot vary the DataShare Open Backup volumes online until Sun service personnel make the required changes to the A7000.
- Be prepared to make additional changes in case the configuration is not compatible with the open system requirements (for example, more devices are needed than you allowed).

Varying the Mainframe Volumes Offline

Vary the mainframe volumes (if any) on the A7000 offline. Alternatively, if the hardware changes are extensive enough to require it, quiesce the system.

Updating the DataShare Open Backup Configuration

Sun service personnel use the information gathered in the planning process to update the DataShare Open Backup configuration while the volumes are offline or the system is quiescent.

Other work such as physical setup and cabling may also be necessary. If a POR is required, the changes on the A7000 can be done while the POR is in progress (but before the IPL).

Restarting the System

If an IPL or POR is required, restart the system when ready. If an IPL was not required, vary the mainframe volumes back online (including the new open systems host volumes) after DataShare Open Backup is activated.

Initializing

The volumes are ready to use without further action. You cannot initialize the open systems host volumes with ICKDSF. The A7000 DataShare Open Backup configuration utilities were used to initialize the open systems host volumes. Write activity to the control structures of the volumes is permanently inhibited.

Using DataShare Open Backup in Backup and Restore Jobs

The DataShare Open Backup open system host volumes must be backed up as part of a whole volume backup strategy. Logistically, you can back up the DataShare Open Backup volume datasets instead of the DataShare Open Backup volumes, but the datasets are only mappings of an underlying volume. Therefore, whether you back up at the dataset level or the volume level, you must treat it as a volume level backup.

If you back up at the dataset level, do not rely on the last modification date in the VTOC as a means to trigger a backup job (as is often done when running incremental dataset backups). The date on these datasets is set to the time DataShare Open Backup was activated and does not change.

When setting up these backup jobs, back up the entire underlying entity as a unit. For example, the 40-Gbyte database in the chapter example is logically a single entity to the host system even though it is implemented as five open volumes and twenty 3390-3 volumes. Backing up only part of it is not useful.



Caution – Attempting to restore anything less than an entire unit can cause loss of data or a system crash.

From a management perspective, track and control the backup of an entire entity as a single unit no matter how many tapes it comprises. Make this task easier by using a clear naming convention.

Quiescing the Client System

This section describes some of the issues involved with the open systems being quiescent during backup. See Chapter 4 for more specific open system client information.

Quiescing the open system entity ensures that:

- All modified data is destaged from the file system or database cache to the A7000
- Data is not changed once the backup begins

In the case of a database, which typically stores its data on raw partitions and manages its own cache, the database must be quiescent or terminated. Your database system may have procedures to destage the cache and transition to read-only mode for the duration of the backup.

Note – Mounting the client volumes read-only during the backup procedure can affect performance because of concurrent access to the A7000 internal cache.

When the file system is mounted on an ordinary volume, unmount the file system to quiesce the volume.

When backing up a boot volume, you cannot unmount the root file system while the host is up. To back up such a volume, shut down the open system.

Optimizing Backup and Restore Performance

To optimize backup and restore performance, maximize concurrency within the constraints of the available resources.

Tape access is often the principal bottleneck in backup and restore activity. The degree of concurrency you attempt to achieve depends on the paths available to your tape subsystem from the backup server. For example, if you have only two parallel channels to a single tape controller, you are more limited in what you can do than if you have several ESCON channels to multiple tape controllers or virtual tape devices.

Break the workload into as many concurrent jobs as there are channel paths to the A7000, and then adjust up or down depending on backup throughput and whether service levels objectives are being met. Use any real-time monitoring tools to analyze job delays due to resource constraints.

If you have four channel paths to the backup host system from the A7000, divide the work into four or five separate jobs to be run concurrently. Consider other workloads running at the same time; particularly those that might use the same channel paths or the tape subsystem. If you can back the volumes up from two systems, try eight concurrent jobs (four from each system).

Consider the length of time needed to back up an entire unit when scheduling the job. Even with a high degree of concurrency, it can take hours to back up the chapter example database of twenty 3390-3 volumes. Typically, a 3390-3 volume can require up to eight tapes in 18-track format or four tapes in 36-track format even when using the Improved Data Recording Capability (IDRC). The database could require as many as 80 tape volumes when using the 3490 tapes with IDRC.

JCL Examples

Mainframe backup and restore strategies are often an integral part of larger, complex automated storage management environments. The simplistic examples that follow use FDR and DFSMSdss to back up a DataShare Open Backup configuration.

The following JCL examples are based on the example UNIX system (TABLE 2-1) used throughout this chapter.

Example 1

FDR is a commonly used tool in many mainframe installations. Here it is used to back up the example database on the open system.

This example illustrates a technique for keeping each of the 9 Gbyte open volumes together as a unit (the database comprises all five open volumes, and one is useless without the other four). Each job backs up the four 3390 volumes that constitute a single open volume, and five similar jobs are run concurrently to back up all five volumes.

Rather than use generation data groups (GDGs), which are difficult to track, the tapes are named using a Julian date in the final dataset name qualifier. This naming convention can be implemented automatically by most automated job scheduling systems. This Julian date qualifier relates all of the backups for the database as a single entity. The dataset name on the tape can be easily related back to the dataset name of the open volume dataset, even though the backup is done by volume name.

To maximize tape usage and minimize load and unload time, the backup datasets are recorded on the tape as four consecutive files on one *logical* tape consisting of several physical volumes.

The esoteric name AUTOCART refers to high-performance ESCON-attached 3490E tape units with automatic cartridge loaders (this is a local esoteric grouping). The default JCL volume count of five volumes is overridden, as it is expected that the number required will be in excess of this. Tape retention policy is managed by the tape library management system and is not coded in the JCL.

The FDR commands indicate that the compression feature of FDR is not used, as IDRC (DCB=TRTCH=COMP) is turned on for the tape controller by default. The DATA=USED parameter ensures that only the used portions of the volumes are backed up. MAXERR=1 ABENDs the job if errors are encountered, ensuring a 100% valid backup.

In the restore job, the MVS catalog identifies the tape volumes needed, and the tape library system generates a pull list allowing the volumes to be stacked in order in a specific automatic cartridge loader at address F84.

CPYVOLID is set to NO because DataShare Open Backup does not allow volume labels to be altered.

CODE EXAMPLE 2-1 Backup With FDR

```
//FDRBKUP1 JOB ...
//*
//* THIS JOB BACKS UP THE 1ST OF 5 OPEN SYSTEMS VOLUMES THAT
//* CONSTITUTES A DATABASE. IT RUNS
//* CONCURRENTLY WITH THE OTHER 4 JOBS (THUS, NO CONCURRENCY
//* WITHIN THE JOB).
//*
//DUMP EXEC PGM=FDR,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=D
//DISK1 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP00
//DISK2 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP01
//DISK3 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP02
//DISK4 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP03
//TAPE1 DD UNIT=AUTOCART,VOL=( , , ,99),LABEL=1,
//      DISP=( ,CATLG,DELETE),
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ1.D1998322
//TAPE2 DD UNIT=AFF=TAPE1,VOL=( , , ,99,REF=*.TAPE1),LABEL=2,
//      DISP=( ,CATLG,DELETE),
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ2.D1998322
//TAPE3 DD UNIT=AFF=TAPE1,VOL=( , , ,99,REF=*.TAPE1),LABEL=3,
//      DISP=( ,CATLG,DELETE),
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ3.D1998322
//TAPE4 DD UNIT=AFF=TAPE1,VOL=( , , ,99,REF=*.TAPE1),LABEL=4,
//      DISP=( ,CATLG,DELETE),
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ4.D1998322
//SYSIN DD *
      DUMP TYPE=FDR,DATA=USED,MAXERR=1
```

CODE EXAMPLE 2-2 Restore With FDR

```
//FDRRSTR1 JOB ...
//*
//* THIS JOB RESTORES THE 1ST OF 5 OPEN SYSTEMS VOLUMES THAT
//* CONSTITUTES A DATABASE. IT RUNS
//* CONCURRENTLY WITH THE OTHER 4 JOBS (THUS, NO CONCURRENCY
//* WITHIN THE JOB).
//*
//DUMP EXEC PGM=FDR,REGION=2M
//SYSPRINT DD SYSOUT=*
//SYSUDUMP DD SYSOUT=D
//DISK1 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP00
//DISK2 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP01
//DISK3 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP02
//DISK4 DD UNIT=3390,DISP=OLD,VOL=SER=BKUP03
//TAPE1 DD DISP=OLD,UNIT=F84,
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ1.D1998322
//TAPE2 DD DISP=OLD,UNIT=AFF=TAPE1,
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ2.D1998322
//TAPE3S DD DISP=OLD,UNIT=AFF=TAPE1,
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ3.D1998322
//TAPE4 DD DISP=OLD,UNIT=AFF=TAPE1,
//      DSN=FDRBKP.OPEN.DBVOL1.SEQ4.D1998322
//SYSIN DD * RESTORE TYPE=FDR,CPYVOLID=NO
```

Example 2

DFSMSDss is commonly used to back up and restore data in MVS/ESA and OS/390.

This example uses DFSMSDss to back up and restore the 4.8 Gbyte open volume of other data from the chapter example.

As in the previous example, the JCL default volume count is overridden. The backup job tape unit is an esoteric group name that refers to a high performance virtual tape system. The virtual tape has been staged to real tape by the time the restore job is run.

These volumes were implemented as open volume datasets on two 3390-3 volumes. To maximize concurrency, the backup and restore are divided into separate jobs for each 3390 volume (both jobs are shown), which are run at the same time.

Because the dataset allocations do not change, the ALLDATA and ALLEXCP keywords are not needed. COMPRESS is not used because hardware compression (IDRC) is turned on at the tape controller. CANCELERROR ABENDs the jobs for any I/O error. On the restore jobs, REPLACE allows the datasets to be restored into the existing

allocation, WRITECHECK ensures that the data is verified as written, and COPYVOLID is omitted because DataShare Open Backup does not allow volume labels to be altered.

CODE EXAMPLE 2-3 Backup With DFSMSdss

```
//DSSBKP1 JOB ...
//STEP001 EXEC PGM=ADRDUSSU,REGION=6M,PARM='UTILMSG=YES'
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSABEND DD SYSOUT=D
//OUTDD1 DD UNIT=VIRTAPE,VOL=(,,,99),LABEL=1,
//          DISP=(,CATLG,DELETE),
//          DSN=BKUP.OPENVOL.SEQ1.D1998322
//INDD1 DD UNIT=3390,DIP=SHR,VOL=SER=BKUP20
//SYSIN DD *
        DUMP -
            INCLUDE(**) -
            INDD(INDD1) -
            OUTDD(OUTDD1) -
            OPTIMIZE(4) -
            CANCELERROR
//DSSBKP2 JOB ...
//STEP001 EXEC PGM=ADRDUSSU,REGION=6M,PARM='UTILMSG=YES'
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSABEND DD SYSOUT=D
//OUTDD2 DD UNIT=VIRTAPE,VOL=(,,,99),LABEL=1,
//          DISP=(,CATLG,DELETE),
//          DSN=BKUP.OPENVOL.SEQ2.D1998322
//INDD2 DD UNIT=3390,DIP=SHR,VOL=SER=BKUP21
//SYSIN DD *
        DUMP -
            INCLUDE(**) -
            INDD(INDD2) -
            OUTDD(OUTDD2) -
            OPTIMIZE(4) -
            CANCELERROR
```

CODE EXAMPLE 2-4 Restore With DFSMSdss

```
//DSSRST1 JOB ...
//STEP001 EXEC PGM=ADRDUSSU,REGION=6M,PARM='UTILMSG=YES'
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSABEND DD SYSOUT=D
//OUTDD1 DD UNIT=3390,DIP=SHR,VOL=SER=BKUP20
//INDD1 DD DSN=BKUP.OPENVOL.SEQ1.D1998322,DISP=SHR
```

CODE EXAMPLE 2-4 Restore With DFSMSdss

```
//SYSIN DD *
  RESTORE -
    INCLUDE(**) -
    INDD(INDD1) -
    OUTDD(OUTDD1) -
    WRITECHECK - REPLACE -
    CANCELERROR
//DSSRST2 JOB ...
//STEP001 EXEC PGM=ADDRSSU,REGION=6M,PARM='UTILMSG=YES'
//SYSPRINT DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//SYSABEND DD SYSOUT=D
//OUTDD2 DD UNIT=3390,DIP=SHR,VOL=SER=BKUP21
//INDD2 DD DSN=BKUP.OPENVOL.SEQ2.D1998322,DISP=SHR
//SYSIN DD *
  RESTORE -
    INCLUDE(**) -
    INDD(INDD2) -
    OUTDD(OUTDD2) -
    WRITECHECK -
    REPLACE -
    CANCELERROR
```

Example 3

This example illustrates how the 1-Gbyte boot volume can be backed up and restored using IEBGENER. This example uses a generation data group (GDG) for the tape dataset name, a convention requiring the use of a model data set control block (DSCB) in the data control block (DCB) parameter, which is set to the correct parameters (BLKSIZE=27648,RECFM=F) with no overrides required.

Note – IEBGENER is not recommended and this example is for informational purposes only.

CODE EXAMPLE 2-5 Backup With IEBGENER

```
//BACKUP JOB ...
//STEP1 EXEC PGM=IEBGENER,TIME=NOLIMIT
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT1 DD DSN=OPENBU.ROOTVOL,DISP=OLD,UNIT=3390,VOL=SER=BKUP22
//SYSUT2 DD DSN=BACKUP.OPNSYS.ROOTVOL(+1),
//          DISP=(,CATLG,DELETE),DCB=BACKUP.DSCB3390,
//          UNIT=TAPE
```

CODE EXAMPLE 2-6 Restore With IEBGENER

```
//RESTORE JOB ...
//STEP1 EXEC PGM=IEBGENER,TIME=NOLIMIT
//SYSPRINT DD SYSOUT=*
//SYSIN DD DUMMY
//SYSUT2 DD DSN=OPENBU.ROOTVOL,DISP=OLD,UNIT=3390,VOL=SER=BKUP22
//SYSUT1 DD DSN=BACKUP.OPNSYS.ROOTVOL(0),
//          DISP=(,KEEP)
```

Automatically Coordinating Processing With the Open System

If a TCP/IP or SNA connection exists between the mainframe and the open systems host, it may be possible to coordinate the backup processes on the mainframe with the open system. For example, a job step using REXEC or RSH before the backup job can unmount the device or quiesce a database. The final step of a job reverses the process.

In the code fragment in CODE EXAMPLE 2-7, the client host is called `opnhost` and the user `mvsbkup` has shell scripts called `quiesce` and `unquiesce` in the `bin` directory. Steps have been added as a wrapper around the backup step to quiesce and unquiesce the client system.

The entire backup process is done in the `BKUP` step. If the process consists of several jobs (all of which must complete), the `QUIESCE` and `RESTART` steps can be packaged as separate jobs and scheduled at the appropriate time.

Alternatively, the open system may signal the mainframe job scheduling system (using file creation or a mainframe console message) that it is ready to be backed up. Some systems (AIX and Windows NT (SNA Server)) support APPC connections that accomplish similar cooperative processing.

CODE EXAMPLE 2-7 Automatically Coordinating Processing With the Open System

```
//BACKUP JOB ...
//QUIESCE EXEC PGM=RSH,REGION=1M,
//          PARM='-l mvsbkup opnhost sh ~/bin/quiesce'
//SYSPRINT DD SYSOUT=*
//SYSTPCD DD DISP=SHR,DSN=MVSBKUP.TCPIP.DATA
//SYSIN DD DUMMY
//*
//BKUP EXEC PGM=ADRDSSU,COND=(0,NE,QUIESCE),
.
.
.
```

CODE EXAMPLE 2-7 Automatically Coordinating Processing With the Open System

```
//*  
//RESTART EXEC PGM=RSH,REGION=1M,  
//          COND=( (0,LT,QUIESCE) , (0,LT,BKUP) ) ,  
//          PARM='-l mvsbkup opnhost sh ~/bin/unquiesce'  
//SYSPRINT DD SYSOUT=*  
//SYSTPCD DD DISP=SHR,DSN=MVSBKUP.TCPIP.DATA  
//SYSIN DD DUMMY
```

Implications for Business Continuation Planning

Open systems backups are restored in the same way they were created: MVS or OS/390 utility programs running against open systems volumes using a properly configured A7000.

To restore open systems at a hot site in the event of a disaster, you must restore the mainframe at the same site and arrange for a properly configured A7000 to be available. You may not be able to recover at your primary hot site in the event of a large regional disaster. Contract with your hot site provider to make an A7000 available at an alternate site or to get an A7000 on short notice.

Sun service personnel must configure the A7000 for DataShare Open Backup. Arrange for this support to be provided as soon as a disaster is declared and in advance when tests are scheduled.

Using the Sun System as the Backup Server

This chapter contains information necessary to plan, set up, and implement DataShare Open Backup using the Sun server to back up the open system volumes stored on the StorEdge A7000.

- Introduction—page 3-2
- Planning—page 3-4
- Setup—page 3-7
- Using DataShare Open Backup—page 3-10

Introduction

DataShare Open Backup makes the open system data stored on the A7000 available to the existing backup and restore tools on your Sun server.

This document does not provide detailed information on any particular backup software product but does provide the information necessary to understand how to backup and restore the open system data stored on the A7000.

These assumptions are made throughout this chapter:

- SCSI Target Emulation (STE) is used for SCSI devices
- Sun Server (HPC 4000 or equivalent) is connected to the A7000 with STE
- Backup and restore product (for example, Sun Enterprise Netbackup) performs the backups and restores from the Sun backup server
- Sun backup server does its backups to tape
- Sun Enterprise Services sets up the phantom headers and tails

FIGURE 3-1 is an example of the connections between the A7000, several open system clients, and a Sun backup server.

Terminology

The following terms are used in this chapter. The Glossary at the end of this manual contains more general terminology.

Term	Description
STE	SCSI Target Emulation. This Sun product allows the A7000 to appear to external hosts as one or more SCSI target disk drives.
Backup Server	Sun HPC 4000 or equivalent Sun Server.
Backup Window	Time available to back up a quantity of data based on the availability requirements of the business.
Client	Open system client machine. That is, an HP, NT, AIX or Sun.
LUN	SCSI Logical Unit Number.




FIGURE 3-1 DataShare Open Backup with a Sun Backup Server

Planning

You must consider the following issues when deciding your backup and restore strategy using the A7000. Some of the setup on the A7000 must be done by Sun Enterprise Services personnel and is not described in this document. However, you should be aware of the entire process.

When planning your strategy for DataShare Open Backup, consider the following:

- Which volumes must be backed up
- Frequency of the backups
- Performance goals for the backup and recovery processes

Use the worksheet at the end of this chapter when planning your backup strategy. The worksheet contains examples of the type of system information that is required.

Capacity

When doing capacity planning, determine the following:

- Quantity of client disk storage to be managed
- Number of disks
- Availability of the disks
- Policies for backing up the data (See Appendix B for the A7000 capacity and limitations.)
- Time period in which a quantity of data must be backed up (also called the Backup Window)

Example – Determining the Bandwidth Necessary to Back Up Data Within the Backup Window

For example, the data center must back up its 200 Gbytes of data during 10 hours of off-peak time. This is 20 Gbytes per hour (200 Gbytes/10 hours), or approximately 6 Mbytes per second. The backup server must be fast enough to read data from the disk volumes being backed up at 6 Mbytes per second or better. The server I/O paths to the target backup disk volumes must have at least 6 Mbytes per second of available bandwidth and at least 6 Mbytes per second of tape bandwidth configured.

Frequency

Because client volumes must be quiescent during the backup, determine the timing and frequency of backups by the availability and duration of off-peak client usage periods. Choosing nonbusiness or off-peak usage hours lessens the chance of adversely affecting client I/O response time. Schedule volume backups daily, weekly, or monthly based on how much data loss you can risk.

Performance

Optimally configure disk volumes of each client across multiple STE ports to uniformly distribute the read I/O load (to the client volumes) across the ports during the backup process. Parallel I/O results in less time required for the backup and recovery of client volumes.

The performance of the tape library significantly affects the time required for the backup and restore processes. Consult your tape library system documentation or vendor to understand the I/O performance bandwidth of your tape library system and how to optimize it.

Recovery

Note – This section applies only when recovering data from a client file system.

The backup server backs up the client volumes as raw data and is not aware of any underlying file system or partitioning (such as that created by the Windows NT FDSK utility). For example, a volume partitioned by FDSK into a primary partition and extended partition with two logical drives appears to Windows NT as volumes C:, D:, and E:. To the backup server, it appears as a single raw disk with no underlying structure.

To recover specific files or directories without backleveling the entire disk, you must restore the backup to a spare device reserved for this purpose. Then use the client operating system facilities to copy the required data from the restored to the target volume.

File-level recovery involves the following:

1. Reserving one or more volumes on the A7000 as scratch volumes (temporary storage).
2. Doing a full volume restore to one of the scratch volumes.

3. When the restore operation is complete, the client system mounts the scratch volume giving the client access to the file system.
4. The client recovers the partition, directories, or files from the scratch volume to a permanent volume.
5. *Unmounting* the scratch volume. This scratch volume may be greater than one disk and must be at least as large as the volume being restored.

Note – To ensure data consistency for restores, volumes with read and write access must be quiescent before you begin the backup procedure. “Unmount” the file systems with read and write access or remount them as read-only if read access is required during the backup process.

Setup

When using DataShare Open Backup, there is a heterogeneous environment in which the A7000 is used for storage by open system clients and backups of the data on the A7000 are performed by a Sun server. The open system client understands only the format of the data on its own partitions, which it accesses using STE. Because the Sun backup server backs up the data in raw format, the format of the client data is not important. The client partition on the A7000 must be encapsulated in a valid Sun disk header and tail enabling the Sun backup server to understand the partition data on the disk. This header and tail are called the STE phantom header and tail. The following figure illustrates phantom header and tail usage.




FIGURE 3-2 Phantom Header and Tail Usage

Sun service personnel create the phantom header and tail on the A7000. You must have an STE SCSI LUN to the Sun backup server, which maps to the open system requiring backup.

Use the Worksheet at the end of this chapter to define the open systems, A7000, and backup partitions. The worksheet contains examples of the type of system information that is required.

The following figure illustrates the connection of two open system clients and the A7000 using two STE connections and the connection of the Sun backup server to the A7000 using only one STE connection.




FIGURE 3-3 Two Open System Clients and the Sun Server Connected to the A7000

Note – Only client data on the A7000 can be backed up using the Sun backup server. Partitions that are on the client are not supported.

▼ To Set Up an A7000 Client Volume for Use With the Sun Backup Server

Using FIGURE 3-3 as an example, the client systems use the SCSI disks provided by the A7000 (D1 and D2).

Use the worksheet at the end of this chapter to ensure the client disks and the Sun backup server STE connections are mapped correctly. The A7000 STE SCSI connected to the Sun backup server must map to these client partitions.

- 1. Determine the open system volumes needed by client. These are D1 and D2 in FIGURE 3-3.**
- 2. Sun service personnel configure the A7000 ensuring the STE SCSI configurations for the clients and to the Sun backup server are correct. The entire client volume must be encapsulated in the STE to the Sun backup server. The phantom header is provided to the Sun backup server. This is where the Sun labels are placed.**
- 3. Because a SCSI channel may have been added, the administrator on the server must modify the `/kernel/drv/sd.conf` file to include the new target and LUNs to support the new devices (refer to the `sd` man page), and then type the following commands:**

```
# touch /reconfigure
# reboot
```

Use the `format` command to set up the slices and labels of the new devices that the A7000 provides. Refer to the `format(1M)` man page. The slices allow the server to read the SCSI devices with the proper offsets to the client data. Each slice or partition can have a tag, as well as a volname for the entire device.

Using DataShare Open Backup

Backing up the data on the A7000 is similar to backing up the SCSI disks on your Sun server. The A7000 is connected to the Sun backup server using an STE SCSI bus connection. Think of the volumes as regular SCSI devices.



Caution – All activity from the open system client to the A7000 *must* be quiescent before beginning the back up procedure. If all volumes being backed up are not quiescent, the backup may be corrupted.

See Chapter 4 for a description of this operation on each supported open system client.

Restoring data is the inverse of backing up data. With DataShare Open Backup, the entire volume is restored. That is, all partitions, datasets, and file systems on the volume are restored to the state at which the backup was performed.

See “Recovery” on page 3-5 for information on restoring files using the scratch volume.

Note – The following procedures are general. See Chapter 4 for more client-specific details.

▼ To Back Up Data

1. **Quiesce all activity on the volumes to be backed up.**
2. **Take these volumes offline (or mount as read-only) on the client machine.**

Note – Mounting the client volumes read-only during the backup procedure can affect performance. Avoid doing this if possible.

3. **Use your backup software package (for example, Netbackup) to back up the data. (Refer to your vendor documentation for detailed information on your backup program.)**
4. **Bring the volumes online.**
5. **Resume activity on the volumes.**

▼ To Restore Data

1. Quiesce all activity on the volumes to be restored.
2. Take these volumes offline on the client machine.
3. Use your restore software package to restore the data.
4. Bring the volumes online.
5. Resume activity on the volumes.

Planning Worksheet

This worksheet contains sample information. The second worksheet is blank.

Open System Name	Client		Open System SCSI BUS Connection and LUN
	Open System Partition or Volume	Open System Size	
NT1	C:	4.0 Gbytes	SCSI bus1,LUN 0

STE SCSI BUS OS Connection and LUN	A7000 Device Name	A7000			STE SCSI Bus Sun Server Connection and LUN
		A7000 Device Type	A7000 Real Device Size	A7000 Phantom header	
STE SCSI bus 1, LUN 0	/dev/rdisk/0d4	MP, VP, Linear, Raid 5	4.1 Gbytes	/dev/rdisk/4d5	STE SCSI bus 2, LUN 0

Sun Backup Server	
SCSI BUS Connection and LUN	Disk device
SCSI BUS 0 LUN 0	/dev/rdisk/c0t0d0s0

Open System Name	Client		Open System SCSI BUS Connection and LUN
	Open System Partition or Volume	Open System Size	

STE SCSI BUS OS Connection and LUN	A7000 Device Name	A7000		A7000 Phantom header	STE SCSI Bus Sun Server Connection and LUN
		A7000 Device Type	A7000 Real Device Size		

Sun Backup Server	
SCSI BUS Connection and LUN	Disk device

Client Issues

The chapter contains information that you must be aware of when using DataShare Open Backup with various platforms as clients.

- Introduction—page 4-2
- Sun Solaris Operating Environment 2.5.1 and 2.6—page 4-7
- Windows NT 4.0—page 4-11
- HP-UX 10.01 and 10.20—page 4-15
- IBM AIX 4.1.2, 4.1.3, and 4.2—page 4-19

Introduction

DataShare Open Backup runs on the StorEdge A7000, which is a storage platform used for data storage and retrieval by other systems (clients). The A7000 sees only data that has been written to the disk, not pending writes held in buffer cache. This is unlike traditional backup programs that run on the same system that uses the data and have automatic access to the disk, as well as to the data that is in buffer cache but not yet written to the disk.




FIGURE 4-1 Relationship Between A7000, Backup Server, and Open System Client

Because the A7000 sees only data that has been written to disk, DataShare Open Backup requires that any volumes being backed up must be in a quiescent state. Quiescence is the state where:

- All data has been written to the disk
- All pending writes in buffer cache have been written to disk
- No write activity occurs during the backup

Quiescence ensures that all of the expected data is backed up. Usually, this requires that the volume be unmounted.

The restore has similar issues. If the volume being restored is not in a quiescent state, pending writes in buffer cache on the open system host overwrite and invalidate what has been restored to the disks residing on the A7000.

Disk Partitions

The traditional backup program deals with local hard disks that have been manipulated and maintained with local tools. Logical Volume Managers or similarly named tool sets can combine many disks into one volume or divide a disk into many volumes.

Using DataShare Open Backup, you can back up only the volumes residing on an A7000. DataShare Open Backup does not see the disks that are locally connected to your open system. The A7000 does not know how your Logical Volume Manager divided or combined the data it presented to your open system as Logical Unit Numbers (LUNs) over SCSI. However, you can set up DataShare Open Backup to back up and restore even complex disk organizations if you know how your disks are organized. Figures 4-2 through 4-6 illustrate various types of disk organizations and show you what to look for when configuring DataShare Open Backup.

The following figure illustrates a single volume comprising three A7000 disks.




FIGURE 4-2 Single Volume

The following figure illustrates three volumes residing on a single A7000 disk.




FIGURE 4-3 Three Volumes on a Disk

The following figure illustrates a single volume comprising two A7000 disks and a locally connected disk (that is, not on the A7000).




FIGURE 4-4 Two Disks in a Single Volume

While Figures 4-2, 4-3, and 4-4 contain valid configurations for data storage and retrieval, DataShare Open Backup can back up only the examples in Figures 4-2 and 4-3 because all of the disks must reside on the A7000. The local disk in Figure 4-4 is not visible to the A7000 and cannot be backed up with DataShare Open Backup.

Figures 4-2, 4-3, and 4-4 also illustrate that there may not be a one to one relationship between the number of volumes on the open system that must be quiescent and the number of entities being backed up on a mainframe or Sun backup server.

Figures 4-5 and 4-6 illustrate that the A7000 data storage presented to an open system looks like a single disk but actually comprises one or more components.

In the following figure, a LUN is composed of a single disk partition, partition 5. (A disk partition may be part of a disk or an entire disk.)




FIGURE 4-5 Data as a Single Disk

In the following figure, a LUN is composed of a Virtual Partition (VP) comprising three disk partitions. VPs are also known as striped partitions or RAID 0.




FIGURE 4-6 LUN Composed of a Virtual Partition

LUNS can comprise mirrored partitions (RAID 1) or other RAID levels.

Figures 4-5 and 4-6 also illustrate that there may not be a one to one relationship between the number of volumes that must be quiescent on the open system and the number of entities being backed up on a mainframe or Sun backup server.

Because configurations can become complex, consult Sun Professional Services for all but the simplest applications.

Device Types

Disks are typically used as blocked or character devices. Blocked devices are file systems that have a structure defined by the computer manufacturer, with named directories and named files. Character (or raw) devices have a structure defined by the user or application. Because the server doing the backup or restore may not be from the same manufacturer that built the open system that uses the data, the server can back up or restore using only raw I/O. Therefore, any backup or restore is always of the full volume type, not incremental.

Changing Disk Hardware Configurations

After changing the disk hardware configuration, the open system must rescan the hardware to load the proper drivers, assign device addresses, etc. Additional steps for partitioning, formatting, and file system generation may be required. Refer to your open system vendor documentation for details.

Restoring Elsewhere

It may be useful to restore a volume somewhere other than where the data was originally backed up.

For example, a backup is performed from `/dev/rhdisk4` and later one file is deleted. Rather than overwriting `/dev/rhdisk4` to restore the one file, you can restore the data on the empty `/dev/rhdisk9`, and then mount `/dev/rhdisk9` temporarily and extract the desired file. In general, UNIX platforms allow this to be done as you would restore and mount any normal volume. Windows NT requires additional steps. Refer to the following sections on specific platforms to determine whether there are any issues for your open system platform.

Sun Solaris Operating Environment 2.5.1 and 2.6

Refer to the *Solaris System Administrator's Guide* or *Sun Volume Manager System Administration Guide* for configuration details for your Solaris operating environment revision.

Disk Naming Conventions

Sun disks have block and raw (character) device files. Each type of device file has its own subdirectory in `/dev`:

- `/dev/dsk` is the block interface
- `/dev/rdisk` is the raw interface

Some commands (for example, `mount`) use the block interface device name from the `/dev/dsk` directory to specify the disk device. Other commands (for example, `newfs`) use the raw interface device name from the `/dev/rdisk` directory to specify the disk device.

These are the disk naming conventions for SCSI disks:

`/dev/dsk/cWtXdYsZ`

or

`/dev/rdisk/cWtXdYsZ`

Where:

`w` is the logical controller number. `W` is always 0 if you have only one SCSI controller on your client system.

`x` is the physical bus target number.

`y` is the drive number. This is also known as the Logical Unit Number (LUN).

`z` is the partition or slice number (0 to 7). To specify the entire disk, use partition 2.

▼ To Mount a File System

- **Bring a disk online by becoming superuser and typing the following:**
 - If you are using a UNIX file system:

```
Sun$ su
Sun# mount /mountpoint
```

- If you are using the Volume Manager:

```
Sun$ su
Sun# vxrecover -s volume_name
```

▼ To Unmount a File System

- **Make the disk quiescent and unavailable to other users by becoming superuser and typing the following:**
 - If you are using a UNIX file system:

```
Sun$ su
Sun# umount /mountpoint
```

Note – The `umount` fails if any users have `/mountpoint` as a part of their current working directory.

- If you are using the Volume Manager:

```
Sun$ su
Sun# vxvol stop volume_name
```

▼ To Backup a Sun System

1. Make the Sun volumes quiescent.

- a. Ensure that there is no write activity during the backup, and that all pending writes in buffer cache are written to disk.
- b. Ensure that there is no current working directory inside the volume, as the volume cannot be unmounted if it contains a working directory.
- c. Ensure that no database, or other application, can access the volume in raw mode.
- d. If the disk is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make the boot disk or volume quiescent, shut down the computer.
- e. If backing up a file system, unmount the volume:

```
Sun# umount /mountpoint      UNIX file systems  
or  
Sun# vxvol stop volume_name  Volume Manager
```

2. Perform the backup from the desired server.

3. Re-enable the Sun volumes.

- a. If the volume was a boot volume, restart the computer.
- b. If backing up database or other application disks, start the database or application.
- c. If backing up a file system, mount the volumes:

```
Sun# mount /mountpoint      UNIX file systems  
or  
Sun# vxrecover -s volume_name  Volume Manager
```

▼ To Restore a Sun System

1. Make the Sun volumes quiescent.

- a. Ensure that there is no write activity during the restore, and that all pending writes in buffer cache are written to disk.

Any pending writes overwrite the data being restored.

- b. Ensure that there is no current working directory inside the volume, as the volume cannot be unmounted if it contains a working directory.

- c. Ensure that no database or other application can access the volume in raw mode.

- d. If the disk in question is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make a boot disk or volume quiescent, shut down the computer.

- e. If restoring a file system, unmount the volume:

```
Sun# umount /mountpoint      UNIX file systems
or
Sun# vxvol stop volume_name  Volume Manager
```

2. Perform the restore from the desired server.

3. Re-enable the Sun volumes.

- a. If the volume was a boot volume, restart the computer.

- b. If restoring a database or other application disks, start the database or application.

- c. If restoring a file system, mount the volumes:

```
Sun# mount /mountpoint      UNIX file systems
or
Sun# vxrecover -s volume_name  Volume Manager
```

Windows NT 4.0

The Microsoft Windows NT Server 4.0 Resource Kit provides detailed information on the Windows NT Server operating system, and includes utilities and tools that help you implement and manage your system.

Online help provides references and how-to-information for all Windows NT tasks.

Disk Naming Conventions

Using Windows NT, you can specify the drive letter for each drive or the drive letters can be assigned automatically. Use the letters C through Z. The letter B is available if the system does not have a second floppy drive. Reserve a drive letter for each CD-ROM drive on the computer.

▼ To Mount a File System

Follow these steps to bring the disk online:

1. **Select Disk Administrator from the Administrative Tools window.**
2. **Select the disk to bring online.**
3. **Select Assign Drive Letter from the Tools menu. This takes effect immediately.**

▼ To Unmount a File System

Follow these steps to make the disk quiescent and unavailable to other users:

1. **Select Disk Administrator from the Administrative Tools window.**
2. **Select the disk that you want to be quiescent.**
3. **Select Drive Letter from the Tools menu.**
4. **Select Do Not Assign Drive Letter.**

▼ To Backup a Windows NT System

1. **Make the Windows NT volumes quiescent.**
 - a. **Ensure that there is no write activity during the backup, and that all pending writes in buffer cache are written to disk.**
 - b. **Ensure that no database or other application can access the volume in raw mode.**
 - c. **If the disk is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make the boot disk or volume quiescent, shut down the computer.**
 - d. **To make a nonboot disk quiescent and unavailable to other users follow these steps:**
 - i. **Select `Disk Administrator` from the `Administrative Tools` window.**
 - ii. **Select the disk that you want to be quiescent.**
 - iii. **Select `Drive Letter` from the `Tools` menu.**
 - iv. **Select `Do Not Assign Drive Letter`.**
2. **Perform the backup from the desired server.**
3. **Re-enable the Windows NT volumes.**

After the backup is complete, resume normal open system operation.
4. **Reboot the Windows NT system if it was shut down. Otherwise, do the following:**
 - a. **Select `Disk Administrator` from the `Administrative Tools` window.**
 - b. **Select the disk to bring online.**
 - c. **Select `Assign Drive Letter` from the `Tools` menu. This takes effect immediately.**

▼ To Restore a Windows NT System

1. **Make the Windows NT volumes quiescent.**
 - a. **Ensure that there is no write activity during the backup, and that all pending writes in buffer cache are written to disk.**
 - b. **Ensure that no database or other application can access the volume in raw mode.**
 - c. **If the disk is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make the boot disk or volume quiescent, shut down the computer.**
 - d. **To make a nonboot disk quiescent and unavailable to other users follow these steps:**
 - i. **a. Select Disk Administrator from the Administrative Tools window.**
 - ii. **Select the disk that you want to be quiescent.**
 - iii. **Select Drive Letter from the Tools menu.**
 - iv. **Select Do Not Assign Drive Letter.**
2. **Perform the restore from the desired server.**
3. **Re-enable the Windows NT volumes.**

After the restore is complete, resume normal open system operation.
4. **Reboot the Windows NT system if it was shut down. Otherwise, do the following:**
 - a. **Select Disk Administrator from the Administrative Tools window.**
 - b. **Select the disk to bring online.**
 - c. **Select Assign Drive Letter from the Tools menu. This takes effect immediately.**

▼ To Write a Signature to the New Disk

If the restore was made to a disk other than the original disk, the Disk Administrator indicates that the restored disk does not have a signature.

1. **Select `Disk Administrator` from the Administrative Tools window.**
A message appears that one or more disks have been added and the system configuration will be updated.

2. **Click on OK.**

A Confirmation window is displayed indicating a signature has not been found on the new disk. Write a signature to the new disk to allow the Disk Administrator to access the drive.

3. **Click on Yes.**

▼ To Label the New Disk

1. **Highlight the new disk.**
2. **Select `Assign Drive Letter` from the Tools menu.**
3. **Type the letter for the new disk.**
4. **Click on OK.**

HP-UX 10.01 and 10.20

Refer to the Hewlett Packard document *Configuring HP-UX for Peripherals* and the system man pages for configuration details.

Disk Naming Conventions

HP-UX disks have block-special and raw (character-special) device files. Each type of device file has its own subdirectory in `/dev`:

- `/dev/dsk` is the block-special interface.
- `/dev/rdsk` is the character-special interface

Commands (for example, `mount`) use the block interface device name from the `/dev/dsk` directory to specify the disk device. Other commands require the raw interface device name from the `/dev/rdsk` directory to specify the disk device.

These are the disk naming conventions for SCSI disks:

```
/dev/dsk/cx[ty]dn[sm]
```

or

```
/dev/rdsk/cx[ty]dn[sm]
```

Where:

x is the logical controller number. This number is a required field that identifies the instance of the interface card.

y is the physical bus target number. This field is optional.

n is the device unit number. This is also known as the Logical Unit Number (LUN). This field is required.

m is the partition or section number. Use partition 0 to specify the entire disk. This field is optional and defaults to 0.

▼ To Mount a File System

- **Bring the disk online by becoming superuser and typing the following:**

```
HP/9000$ su
HP/9000# mount /mountpoint
```

▼ To Unmount a File System

- **Make the disk quiescent and unavailable to other users by becoming superuser and typing the following:**

```
HP/9000$ su
HP/9000# umount /mountpoint
```

Note – The `umount` fails if any users have `/mountpoint` as a part of their current working directory.

▼ To Backup an HP-UX System

1. **Make the HP volumes quiescent.**
 - a. **Ensure that there is no write activity during the backup, and that all pending writes in buffer cache are written to disk.**
 - b. **Ensure that there is no current working directory inside the volume, as the volume cannot be unmounted if it contains a working directory.**
 - c. **Ensure that no database, or other application, can access the volume in raw mode.**
 - d. **If the disk is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make the boot disk or volume quiescent, shut down the computer.**
 - e. **If backing up a file system, unmount the volume:**

```
HP/9000# umount /mountpoint
```


c. If restoring a file system, mount the volumes:

```
HP/9000# mount /mountpoint
```

Note – You can restore a volume to a location that is different from where the backup was done.

IBM AIX 4.1.2, 4.1.3, and 4.2

Refer to the *IBM AIX System Management Guide: Operating System and Devices* for your revision level of AIX, and the system man pages, for configuration details.

Disk Naming Conventions

AIX disks have block-special and raw (character-special) device files. Each type of device file has its own subdirectory in `/dev`:

- `/dev/hdisk` is the block-special interface
- `/dev/rhdsk` is the raw-special interface

Some commands (for example, `mount`) use the block interface device name from the `/dev/hdisk` directory to specify the disk device. Other commands use the raw interface device name from the `/dev/rhdsk` directory to specify the disk device.

These are the disk naming conventions for SCSI disks:

```
/dev/hdiskNN
```

or

```
/dev/rhdskNN
```

Where:

`NN` is the disk number assigned by system. This number has no relationship to the hardware address of the disk but is the next available number.

▼ To Mount a File System

- **Bring the disk online by becoming superuser and mounting file systems with entries in `/etc/fstab`:**

```
IBM$ su
IBM# mount /mountpoint
```

▼ To Unmount a File System

- **Make the disk quiescent and unavailable to other users by becoming superuser and typing the following command:**

```
IBM$ su
IBM# umount /mountpoint
```

Note – The `umount` fails if any users have `/mountpoint` as a part of their current working directory.

▼ To Backup an IBM AIX System

1. **Make the AIX volumes quiescent.**
 - a. **Ensure that there is no write activity during the backup, and that all pending writes in buffer cache are written to disk.**
 - b. **Ensure that there is no current working directory inside the volume, as the volume cannot be unmounted if it contains a working directory.**
 - c. **Ensure that no database or other application can access the volume in raw mode.**
 - d. **If the disk is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make the boot disk or volume quiescent, shut down the computer.**
 - e. **If backing up a file system, unmount the volume:**

```
IBM# umount /mountpoint
```

2. **Perform the backup from the desired server.**
3. **Re-enable the AIX volumes.**
 - a. **If the volume was a boot volume, restart the computer.**
 - b. **If backing up database or other application disks, start the database or application.**

c. If backing up a file system, mount the volumes:

```
IBM# mount /mountpoint
```

▼ To Restore an IBM AIX System

1. Make the AIX volumes quiescent.

- a. Ensure that there is no write activity during the restore, and that all pending writes in buffer cache are written to disk. Any pending writes overwrite the data being restored.
- b. Ensure that there is no current working directory inside the volume, as the volume cannot be unmounted if it contains a working directory.
- c. Ensure that no database, or other application, can access the volume in raw mode.
- d. If the disk in question is a boot disk or boot partition, the system will not run with the boot partition unmounted. To make a boot disk or volume quiescent, shut down the computer.
- e. If restoring a file system, unmount the volume:

```
IBM# umount /mountpoint
```

2. Perform the restore from the desired server.

3. Re-enable the AIX volumes.

- a. If the volume was a boot volume, restart the computer.
- b. If restoring database or other application disks, start the database or application.
- c. If restoring a file system, mount the volumes:

```
IBM# mount /mountpoint
```

Note – You can restore a volume to a location that is different from where the backup was done. The AIX disk naming convention of using the next available number as opposed to having a device hardware location based name makes locating the new disk difficult. Inventory existing disks before adding new ones.

Best Business Practices

This chapter contains best business practices for specific configurations, which will help you understand how to plan and set up DataShare Open Backup.

- Introduction—page 5-2
- Mainframe Backup of Open Systems Data—page 5-2
- Sun Backup of Open Systems Data—page 5-7

Introduction

This chapter contains examples of how to use a mainframe or a Sun system as a backup server to back up your open systems data. Each example contains the following information:

- Current backup and restore methods
- Hardware and configurations
- Analysis of the problem
- Proposed method for using DataShare Open Backup to solve the problem
- Backup and restore considerations

Mainframe Backup of Open Systems Data

A mainframe site is upgrading a Sun StorEdge A7000 for additional capacity. Until now, the A7000 has been used only for mainframe DASD. The additional capacity will host data for a Windows NT server running a sales force support application and a Sun server running a financial application.

The open systems volumes being added to the A7000 must be backed up.

Current Method

This section describes the current mainframe and open systems environments.

Mainframe Environment

The mainframe uses Fast Dump Restore (FDR) to do full volume backups of most of DASD on a staggered weekly or monthly basis depending on volatility.

Backups of mainframe volumes using Automated Backup and Restore (ABR) are scheduled daily. These jobs are scheduled automatically based on resource availability, time, and system events with CA-7. They are backed up to a 3490E-compatible robotic tape system using IDRC and 36-track (bidirectional) format. The CA-1 tape management system controls tape usage, retention, and reporting. The mainframe system is reliable and automated, requiring minimal staffing.

Open Systems Environment

Current backup and restore on open systems vary by platform,. Some automation exists for open systems, but it varies and requires more staff intervention than for the mainframe. Backup and restore on open systems is less reliable than on the mainframe. Experience suggests that media are often undependable and comparatively fragile, resulting in unreadable tapes.

Application Details

The mainframe has near 24x7 availability requirements.

The open systems must be available during a business day that spans several time zones. Downtime can be scheduled during off-hours.

The interface to the A7000 for the mainframe is ESCON and mixed BMC, ESCON elsewhere.

The interface to the A7000 for the open systems is an Ultra-SCSI.

Mainframe Hardware and Configuration

- IBM 9021
- OS/390
- TCP/IP

Open Systems Hardware and Configuration

- A Compaq Prolinea dual Pentium Pro NT server with:
 - Windows NT server 4.0
 - A7000 providing 18-Gbytes RAID 5 storage plus 2-Gbytes RAID 1 for the boot device
- A Sun HPC 4000 server with:
 - Sun Solaris Operating Environment 2.5
 - A7000 provides 32-Gbytes RAID 5 storage plus 2-Gbytes RAID 1 for the boot device

Problem Analysis

- Full volume backup required
- A server must be shut down when backing up its boot volumes
- Nonboot volumes may be backed up when quiescent

Solution Description

DataShare Open Backup backs up open systems volumes on the A7000 using existing mainframe backup facilities. The open system volumes appear to the mainframe backup software as single large datasets on each volume.

Proposed Method of Accomplishing Solution

The proposed method includes

- Capacity planning
- Configuring the A7000 (Completed by Sun Enterprise Services Personnel)
- Configuring the mainframe

▼ To Do Capacity Planning

- 1. Define mainframe volumes as 3390-3 volumes, as required by the existing I/O configuration.**
- 2. A7000 DataShare Open Backup half-track blocking provides 2.58 Gbytes of storage per 3390-3 device. Calculate the amount of data to be backed up from open systems storage and determine the number of mainframe volumes needed to support it.**

Windows	18-Gbytes (two 9-Gbyte volumes)	=	Eight 3390-3 volumes
NT:	2-Gbyte C: Drive	=	One 3390-3 volume
Sun:	32-Gbytes (four 8-Gbyte volumes)	=	Sixteen 3390-3 volumes
	2-Gbyte boot device *	=	One 3390-3 volume
Total:	54-Gbytes		26 volumes

* The mirror of the RAID 1 boot device does not need to be backed up.

▼ To Configure the A7000 (Completed by Sun Enterprise Services Personnel)

- **Map the open systems volumes defined in the preceding table to the respective mainframe 3390 volumes and datasets.**

(The data center administrator selects the following names consistent with the installation's naming convention of BU (Back Up) and the two character open system identifier NT (Windows NT) or SN (Sun).)

Volume serial names:	BUNT00-BUNT08	BUSN00-BUSN16
Dataset name:	BU.NTVOL01.SEQ1	BU.SUNVOL01.SEQ1
	BU.NTVOL01.SEQ2	BU.SUNVOL01.SEQ2
	BU.NTVOL01.SEQ3	BU.SUNVOL01.SEQ3
	BU.NTVOL01.SEQ4	BU.SUNVOL01.SEQ4
	BU.NTVOL02.SEQ1	BU.SUNVOL02.SEQ1
	.	.
	.	.
	.	.
	BU.NTVOL02.SEQ4	BU.SUNVOL04.SEQ4

▼ To Configure the Mainframe

- **Map the volumes defined in the preceding table to the mainframe hardware. Assign CHPID 2E and 1E and device addresses 220-23F for this function using the interactive ISPF Hardware Configuration and Definition dialog.**

Note – Two CHPIDs and 32 device addresses are assigned to this function to reserve them for future expansion. Only the 26 addresses, 220 thorough 239, are used at this time.

▼ To Configure the Open Systems

There are no special configuration issues. These volumes are configured normally to the respective systems.

Application Solution

FDR DUMP jobs are created to back up the volumes nightly. To optimize concurrency, an average of four volumes is backed up by each job in a single step.

Backup Considerations

- Backup jobs are added to the nightly CA-7 production schedule on the mainframe.
- Windows NT backups are scheduled to run at midnight (when the volumes are quiescent).
- Sun backup jobs take advantage of the TCP/IP network connection to the Sun machine. An event (a mainframe operator console message) triggers the automatic start of the backup job. The event is created by a shell script on the Sun machine, which quiesces the database and starts an REXX exec on MVS using TCP/IP. When the backup is complete, CA-7 runs another job that restarts the database on the Sun machine using the MVS RSH program.
- Backups of the Sun and Windows NT boot packs are done on a monthly basis during scheduled maintenance.

Restore Considerations

After the problem that caused the outage is repaired, the volumes are varied online to the mainframe, and the FDR restore jobs are submitted interactively using the FDR ISPF panels to automatically select the latest backup.

Sun Backup of Open Systems Data

A Sun site is upgrading a Sun StorEdge A7000 for additional capacity. Until now, the A7000 has been used only for Sun storage. The additional capacity will host data for a Windows NT server running a sales force support application and a Sun server running a financial application.

The open systems volumes being added to the A7000 must be backed up.

Current Method

The Sun server uses NetBackup to do full volume backups of all storage on a staggered weekly or monthly basis depending on volatility.

Backups using NetBackup are scheduled daily. Database disks are backed up using the database utilities. Schedules are set using the NetBackup management facilities. They are backed up to a 400-Gbyte/8mm autoloading tape system.

Backup and restore on open systems vary by platform.

Application Details

The open systems must be available during a business day that spans several time zones. Downtime can be scheduled during off-hours.

The interface to the A7000 for the Sun server and the open systems is SCSI-II fast and wide.

Sun Hardware and Operating System

- 10-processor E6000
- Sun Solaris operating environment 2.6
- A7000 configured as SCSI drives of various sizes

Open Systems Hardware and Operating System

- A Compaq Prolinea dual Pentium Pro NT server with the following:
 - Windows NT server 4.0
 - A7000 providing 18-Gbytes RAID 5 storage plus 2-Gbytes RAID 1 for the boot device
- A Sun HPC 4000 server with the following:
 - Sun Solaris operating environment 2.5.1
 - A7000 providing 32-Gbytes RAID 5 storage plus 2-Gbytes RAID 1 for the boot device

Problem Analysis

- No show stoppers are apparent. Hardware and software environments match supported configurations.
- Windows NT server can be shut down during the backup operation and HPC 4000 volumes can be “unmounted.”
- Full volume backups of the open systems are required.

Solution Description

DataShare Open Backup backs up open systems volumes on the A7000 using existing backup facilities. The open systems volumes appear to the E6000 backup software as raw devices.

Capacity Planning

When doing capacity planning, determine the following:

- Amount of client disk storage to be managed
- Number of disks
- Availability of the disks
- Policies for backing up the data
- Time period in which a quantity of data must be backed up (also called the Backup Window)

Configuring the Open Systems

There are no special configuration issues. These volumes are configured normally to the respective systems.

Application Solution

NetBackup jobs are created to back up the volumes nightly. To optimize concurrency, an average of four volumes is backed up by each job in a single step.

Backup Considerations

- The Windows NT backups are scheduled to run at midnight (when the volumes are quiescent).
- The Sun backup job takes advantage of the TCP/IP connection with the Sun machine. An event triggers the automatic start of the backup job. The event is created by a shell script on the Sun machine, which quiesces the database. When the backup is complete, NetBackup runs another job that restarts the database on the Sun machine using the remote execution.
- Backups of the Sun and Windows NT boot packs are done on a monthly basis during scheduled maintenance.

Restore Considerations

After the problem that caused the outage is repaired, the devices are mounted on the E6000 and the NetBackup restore jobs are submitted interactively to automatically select the latest backup.

A7000 Platform-Specific Storage Information

This appendix provides basic information and procedures that enable you to set up your supported open system platforms for use with the A7000. Refer to your vendor documentation for details on configuring and using your system.

- Sun Solaris Operating Environment 2.5.1 and 2.6—page A-2
- Windows NT 4.0—page A-6
- HP-UX 10.01 and 10.20—page A-10
- IBM AIX 4.1.2, 4.1.3, and 4.2—page A-18

Sun Solaris Operating Environment 2.5.1 and 2.6

Refer to the *Solaris System Administrator's Guide* or *Sun Volume Manager System Administration Guide* for configuration details for your Solaris operating environment revision.

▼ To Scan or Rescan the Hardware

There are two methods to scan the hardware and set up the associated device drivers and nodes:

- Use the `boot -r` command when the system is shut down and waiting to be booted:

```
Sun# boot -r
```

- Use the DR Interface command sequence when the system is booted.

```
Sun# drvconfig; devlinks; disks;
```

▼ To Format and Partition Disks

1. Become superuser.

```
Sun$ su
```

2. Type `format` and press the `Enter` key. A list of available disks is displayed.
3. Type the number of the new disk. (The number is displayed in the list.) The `Format` menu and the `format>` prompt are then displayed.
4. Type the partition and press the `Enter` key.
5. Select the appropriate option from the `Partition` Menu and partition the disk.
6. When the partitioning is complete, type `label` and press the `Enter` key.

7. Type `quit` and you are returned to the `format` prompt.
8. Type `quit` and press the `Enter` key.

▼ To Verify the Disk Partitioning

- Type this command and press the `Enter` key. Information for the disk you specified is displayed.

```
Sun# prtvtoc /dev/rdisk/cWtXdYs2
```

▼ To Make a UNIX File System (UFS)

A disk must be formatted, partitioned, and labeled before you can create a UFS on it.

1. Become superuser.

```
Sun$ su
```

2. Type the following, and you are asked if you want to proceed.

```
Sun# newfs /dev/rdisk/cWtXdYsZ
```

3. Check the file system.

```
Sun# fsck /dev/rdisk/cWtXdYsZ
```

Mounting File Systems and Starting Volumes

Like mounting a file system, starting a volume affects its user-availability. Starting a volume changes its state and makes it available.

Starting a volume changes the volume state from `DISABLED` or `DETACHED` to `ENABLED`. The success of this operation depends on the ability to enable a volume. If a volume cannot be enabled, it remains in its current state.

▼ To Use the Volume Manager (VxVM)

- To start the volume, type:

```
Sun# vxrecover -s volume_name
```

▼ To Use a UNIX File System

1. Become superuser.

```
Sun$ su
```

2. If the mount point is not created, create the directory.

A mount point is a directory to which the mounted file system is attached.

```
Sun# mkdir /mountpoint
```

3. Mount the file system.

```
Sun# mount /dev/dsk/cWtXdYsZ /mountpoint
```

If the mount point entry is in `/etc/vfstab`, type:

```
Sun# mount /mountpoint
```

4. If this partition is to be mounted or unmounted on reboot or shutdown, edit the `/etc/vfstab` file and add the entry. This is an example of an added `ufs` mount point:

```
#device          device      mount      FS    fsck  mount  mount
#to mount        to fsck    point     type  pass  at boot options
#
/dev/dsk/c0t0d0s0 /dev/rdsk/c0t0d0s0 /usr      ufs   1     yes    -
```

Use the `mountall` command if the `/etc/vfstab` file includes all the entries to be mounted. `mountall` mounts all the mount points configured in `/etc/vfstab` if they are not mounted already.

Unmounting File Systems and Stopping Volumes

Like unmounting a file system, stopping a volume affects its availability to the user. Stopping a volume makes it unavailable.

Stopping a volume changes the volume state from ENABLED or DETACHED to DISABLED. If the command cannot stop the volume, it remains in its current state.

▼ To Stop the Volume Manager (VxVM)

- To stop the volume, type:

```
Sun# vxvol stop volume_name
```

▼ To Unmount a UNIX File System

1. Become superuser.

```
Sun$ su
```

2. Unmount the file system.

```
Sun# mount /mountpoint
```

Note – The `umount` fails if any users have `/mountpoint` as a part of their current working directory.

Windows NT 4.0

The Microsoft Windows NT Server 4.0 Resource Kit provides detailed information on the Windows NT Server operating system, and includes utilities and tools that help you implement and manage your system.

Online help provides references and how-to-information for all Windows NT tasks.

▼ To Obtain SCSI IDs

- 1. Use the Adaptec BIOS during system boot to determine the SCSI host adapter (for example, Adaptec AHA-2944W) and the Sun SCSI Target Emulation (STE) IDs. The Adaptec BIOS displays a message similar to the following during system boot:**

```
Adaptec AHA-2940/2940W/2944W BIOS v1.0.0
(c) 1993 Adaptec, Inc. All Rights Reserved.
< < < Press <Ctrl><A> for SCSISelect (TM) Utility! > > >
```

- 2. Press the Ctrl-A. An options box containing two choices is displayed.**

```
Configure/View Host Adapter Settings
SCSI Disk Utilities
```

3. Use the arrow keys to highlight `SCSI Disk Utilities`, and then press the `Enter` key. The SCSI bus is scanned and the IDs are listed.

```
SCSI ID#  0: No Device
.
.
.
SCSI ID#  5: SUN STE
SCSI ID#  6: No Device
SCSI ID#  7: AHA-2940/2940W/2944W
SCSI ID#  8: No Device
.
.
.
SCSI ID# 14: No Device
SCSI ID# 15: No Device
```

4. Press the `Escape` key twice to display the exit menu.
5. Use the arrow keys to highlight `Yes`, and then press the `Enter` key. The system displays a message to press any key to reboot.

Displaying the Configured Disks

While the Windows NT system is running, you can display all configured disks by double clicking on `My Computer` or by using the `Disk Administrator`.

▼ To Login in as the Disk Administrator

1. Select the `Start` button in the lower left corner of the screen.
2. Highlight `Programs` in the pop up menu.
3. Highlight `Administrative Tools` in the pop up menu.
4. Select `Disk Administrator` in the pop up menu.

This screen takes a few seconds to complete and displays all the disks configured and the type of file systems on the disks.

Note – You can change an existing drive without rebooting the system, but you must reboot the system to add a drive.

Creating a Windows NT File System

The file system is created by formatting and partitioning the disk.

You must log in as the Disk Administrator (refer to the previous section) to format and partition the disk.

After selecting Disk Administrator in the pop up menu, a message similar to the following is displayed. Click on the OK button to initialize the Disk Administrator.

Disk Administrator has determined that one or more disks have been removed from your computer since Disk Administrator was last run, or that one or more disks are off-line. Configuration information about the missing disk(s) will be retained.

▼ To Create a Primary Partition in the Disk Administrator

1. **Select an area marked disk free space.**
2. **Select Create from the Partition menu.**
3. **Type the desired size in the Create partition of Size box.**
4. **Click on OK.**
5. **Select Commit Changes Now from the Partition menu. After the partition is committed, it is marked as Unformatted.**
6. **Answer Yes to Update Configuration File.**
7. **Click on OK to update and create a new Emergency Repair Disk.**
The disk is marked as unknown. Format the disk now.

▼ To Format the Partition

1. **Select the partition in the Disk Administrator window.**
2. **Choose Format from the Tools menu.**
3. **Select the file system type (NTFS or FAT) in the File System box.**
The two available file system types:

- New Technology File System (NTFS) performs better with larger volume sizes (>1.0 Gbytes) and provides superior security and recovery features.
 - File Allocation Table (FAT) is the MSDOS style file system and is compatible with, and allows access to, DOS files. FAT performs best with smaller volumes. Performance degrades as the volume size exceeds 1.0 Gbyte.
4. **Enter a volume label (if desired) in the Label box.**
 5. **Click in the Quick Format check box to format without checking for errors.**
 6. **Click on Start. As the partition is formatted, progress is displayed on a bar graph.**
 7. **Click on OK when the Format Complete message is displayed.**
File system and partition size are automatically updated and are displayed in the new disk area.

HP-UX 10.01 and 10.20

Refer to the Hewlett Packard document *Configuring HP-UX for Peripherals* and the system's man pages for configuration details.

▼ To Scan or Rescan the Hardware

1. Use the HP/UX `ioscan` command to list devices and their associated `/dev` nodes.

```
HP/9000# ioscan -fn -d disc3
```

The following information is returned for each disk. You can ignore the device addresses for floppy disks.

```
Class      I  H/W Path  Driver S/W State H/W Type  Description
=====
disk       10 16.4.0   disc3  CLAIMED  DEVICE  SUN  STE
          /dev/dsk/c0t4d0   /dev/rdisk/c0t4d0
          /dev/floppy/c0t4d0 /dev/rfloppy/c0t4d0
```

If your output does not display `/dev` node information for all of your disks, disks have been added since the last time device entries were built.

2. To build new device entries, type:

```
HP/9000# insf -e
```

Rerunning `ioscan` displays the `/dev` information.

3. To display information about a specific disk, type:

```
HP/9000# /etc/diskinfo /dev/rdisk/c0t4d0
SCSI describe of /dev/rdisk/c0t4d0:
      vendor: SUN
      product id: STE
      type: direct access
      size: 2007180 Kbytes
      bytes per sector: 512
```

Creating an HP-UX File System

You can format and partition disks and create file systems in two ways:

- Using the System Administration Manager
- From a command line

▼ To Format and Partition the Disk and Create a File System Using the System Administration Manager (SAM)

1. Start SAM.
2. Create a volume group, logical volume, and file system for the Sun STE vdisks.
 - a. Name the volume group `vg_sun01`, and have it consist of the SUN STE DISK DRIVE at hardware address 24, SCSI ID 4, LUN 0 (`c2t4d0`). To accomplish this, select the following SAM choices:

```
Disks and File Systems →
Volume Groups →
Actions →
Create or Extend →
24.4.0 SUN STE DISK DRIVE →
OK →
Create or Extend a Volume Group →
Volume Group Name... vgsun01 →
OK
```

- b. Within this volume group, create a 252 Mbyte logical volume named `lvsun10`, which contains the file system that will be mounted at `/sun01`. To do so, select the following:

```
Add New Logical Volumes... →
LV Name: lvsun10 →
LV Size (Mbytes): 252 →
Mount Directory: /sun01 →
Add →
OK →
OK →
Yes
```

- 3. Extend the Volume Group to include the SUN STE DISK DRIVE at hardware address 24, SCSI ID 4, LUN 1 (c2t4d1) by selecting the following:**

```
Disks and File Systems →
Volume Groups →
Actions →
Create or Extend →
24.4.1 SUN STE DISK DRIVE →
OK →
Create or Extend a Volume Group →
Volume Group Name... vgsun01 →
OK →
OK →
Yes
```

- 4. To examine the volume group, select the following:**

```
Disks and File Systems →
Volume Groups →
vgsun01 →
Actions →
Zoom →
(Shows statistics of volume group) →
OK →
View More Information →
(Shows physical volumes)
```

- 5. To examine the logical volume, select the following:**

```
Disks and File Systems →
Logical Volumes →
lvsun10 →
Actions →
View More Information →
(Shows statistics and physical vol)
```

6. To examine the file system, select the following:

```
Disks and File Systems →  
File Systems →  
/sun01 →  
Actions →  
View More Information →  
(Shows statistics and physical vol)
```

▼ To Format and Partition the Disk and Create a File System From the Command Line (Without SAM)

- 1. Create a volume group for the Sun disks. Name the volume group `vgsun01` and have it consist of the SUN STE DISK DRIVE at hardware address 16, SCSI ID 1, LUN 0 (`c0t1d0`).**
 - a. List all existing A7000 volume groups and logical volumes and ensure that any previous failures of the A7000 `vg` and `lv` have been removed:**

```
HP/9000# ls /dev/vgsun*  
HP/9000# ls /dev/lvsun*  
HP/9000# rm -rf /dev/vgsun*  
HP/9000# rm -rf /dev/lvsun*
```

- b. Allocate the device so it can be assigned to a volume group:**

```
HP/9000# pvcreate /dev/rdisk/c0t1d0  
Physical volume "/dev/rdisk/c0t1d0" has been successfully created.  
HP/9000# pvdisplay /dev/rdisk/c0t1d0
```

- c. Create a directory for the volume group:**

```
HP/9000# mkdir /dev/vgsun01
```

d. Create a control file named `group` in the Volume Group Directory:

```
HP/9000# mknod /dev/vgsun01/group c 64 0x010000
```

Where `c` specifies that `group` is a character device file, `64` is the major number for the LVM pseudo driver, and `0x0y0000` is the unique Volume Group Identifier (where `y` is a value, 0 through 9, that must be unique across all volume groups).

Note – The HP documentation describes the final parameter as `0xab0000`, where `ab` is the hexadecimal representation of `nn` in `/dev/vgnn`.

```
HP/9000# ls -l /dev/vgsun01/group
crw-rw-rw- 1 root sys 64 0x010000 May  3 14:49 /dev/vgsun01/group
```

e. Create the Volume Group and assign the A7000 device:

```
HP/9000# vgcreate /dev/vgsun01 /dev/dsk/c0t1d0
```

This updates `/etc/lvmpvg`, an ASCII file containing physical volume group mapping.

f. Display detailed information about the volume group:

```
HP/9000# vgdisplay -v vgsun01
--- Volume groups ---
VG Name                /dev/vgsun01
VG Write Access        read/write
VG Status              available
Max LV                 255
Cur LV                0
Open LV               0
Max PV                 16
Cur PV                1
Act PV                1
Max PE per PV         2045
VGDA                   2
PE Size (Mbytes)      4
Total PE              2045
Alloc PE               0
Free PE               2045
Total PVG             0
--- Physical volumes ---
PV Name                /dev/dsk/c0t1d0
```

PV Status	available
Total PE	2045
Free PE	2045

where Max PE per PV is what is currently available (that is, the original size less overhead) and PE Size is the default physical extent size.

2. Within this volume group, create a logical volume named lvsun10.

a. Create a 2000-Mbyte logical volume named lvsun10:

```
HP/9000# lvcreate -L 2000 -n lvsun10 vgsun01
Logical volume "/dev/vgsun01/lvsun10" has been successfully
created with character device "/dev/vgsun01/rlvsun10".
Logical volume "/dev/vgsun01/lvsun10" has been successfully
extended.
```

b. You can dynamically extend the size of an existing logical volume. Here, the size is extended to 4000 Mbytes:

```
HP/9000# lvextend -L 4000 /dev/vgsun01/lvsun10
HP/9000# lvdisplay -v /dev/vgsun01/lvsun10 | pg
--- Logical volumes ---
LV Name                /dev/vgsun01/lvsun10
VG Name                /dev/vgsun01
LV Permission          read/write
LV Status              available/syncd
Mirror copies          0
Consistency Recovery   MWC
Schedule               parallel
LV Size (Mbytes)       4000
Current LE             500
Allocated PE           500
Stripes                0
Stripe Size (Kbytes)   0
Bad block              off
Allocation              strict/contiguous

--- Distribution of logical volume ---
PV Name                LE on PV  PE on PV
/dev/dsk/c0t1d0        500       500
--- Logical extents ---
LE   PV1                PE1  Status 1
0000 /dev/dsk/c0t1d0    1500 current
...
```

3. **Format the disk. This sets up info in the disk partition table and scans for bad disk blocks.**

```
HP/9000# mediainit /dev/rdisk/c0t1d0
```

4. **Create a mount directory for each raw A7000 volume.**

```
HP/9000# mkdir /sun01
```

5. **Create a file system for each desired A7000 device.**

```
HP/9000# newfs /dev/rdisk/c0t1d0
```

or, when using volume groups:

```
HP/9000# newfs /dev/vgsun01/rlvsun10
```

Options to the `newfs` command include:

<code>-i bytes</code>	The number of bytes per inode (default is 6144).
<code>-F type</code>	The type of file system to create.

6. **When the file system is created, run `fsck` to verify it. For example:**

```
HP/9000# fsck -y /dev/rdisk/c0t1d0
```

7. **Mount the desired file systems.**

```
HP/9000# mount /dev/vgsun01/lvsun10 /sun01
```

To check the number of free blocks and free i-nodes in mounted file systems, use the `bdf` command (a Berkeley variation of `df`):

```
HP/9000# bdf
File system          kbytes    used    avail capacity Mounted on
/dev/vg00/lvol1      99738     48629   41135    54%    /
/dev/vg00/lvol3      141914   109983  17739    86%    /usr
/dev/vgsun01/lvsun10 2017244    10     1815508 0%     /sun01
```

IBM AIX 4.1.2, 4.1.3, and 4.2

Refer to the *IBM AIX System Management Guide: Operating System and Devices* for your revision level of AIX, and the system's man pages for configuration details.

▼ To Scan or Rescan the Hardware

1. Before adding Sun STE disks on the RS6000, type the following command to take an inventory of the existing disk configuration.

```
IBM# lscfg
```

2. You must power off the RS6000 to add SCSI boards or any other hardware modifications. You do not have to power off the machine to change the disk configuration from the A7000. After configuring disks for the RS6000, type the following command to rescan the hardware and perform /dev node assignments:

```
IBM# cfgmgr -v
```

3. List the device configuration to determine the names assigned to the newly configured disks.

```
IBM# lscfg
```

Creating the IBM AIX File System

The System Management Interface Tool (SMIT) is used in the procedures in this section.

▼ To Optimize Device Attributes

Each SCSI device is assigned a default set of attributes that usually need to be adjusted to optimize performance. The most significant of these are:

- Queue depth
- Queue type
- Read/write time out
- Start time out
- Reassign time out

1. You can use SMIT to modify the attributes stored in the system device database by selecting the following menu options:

```
Devices → Fixed Disk → Change/Show Characteristics of a Disk
```

2. Use the arrow keys to select the desired device (for example, `hdisk1`), and then press the `Enter` key.
3. Use the arrow keys to move to the attribute fields and modify the values as follows:

```
Queue DEPTH = 16
Queue TYPE = simple
READ/WRITE time out = 120
START time out = 120
REASSIGN time out = 120
```

Note – Depending on the system load, you may need to set the `READ/WRITE` time out value higher.

When you are ready to commit your changes, press the `Enter` key.

4. Repeat Steps 2 and 3 for the remaining disks.

▼ To Use the SMIT to Configure Sun Disks

1. Choose a volume group for the Sun disks.

- a. To create a new volume group, start SMIT, and then navigate to the Add Volume Group screen by highlighting and selecting the following menu items:

```
IBM# smit
System Storage Management → Logical Volume Manager → Volume
Groups → Add a Volume Group
```

- b. Type in the desired volume group name (for example, `vgsun01`).
- c. Use the arrow keys to advance to the next entry field, physical partition size. The default value is 4 Mbytes. Pressing the Tab key increments the value by powers of two.
- d. Select the desired physical volume names. Pressing Esc-4 provides a list. Arrow down to `hdisk1` and press Enter.
- e. Accept the defaults for the remaining values and press Enter to create the volume group.
- f. Repeat this process to create volume group `vgsun02` containing physical volume `hdisk2`.

2. Verify the new volume groups.

- a. Press Esc-3 until you are at the Volume Groups screen.
- b. Select the List Contents of a Volume Group function.
- c. Press Esc-4 to list available volume group names. Highlight the desired volume group and press Enter.
- d. Press Enter a second time to display a status listing for the volume group. For example, the status listing for `vgsun02` shows the following:

```
=====
                        COMMAND STATUS
Command: OK                stdout: yes                stderr: no

Before command completion, additional instructions may appear below.

VOLUME GROUP:    vgsun02  VG IDENTIFIER:  00008886b38acfb
VG STATE:        active   PP SIZE:        16 megabyte(s)
VG PERMISSION:   read/write  TOTAL PPs:     255 (4080 megabytes)
MAX LVs:         256       FREE PPs:      255 (4080 megabytes)
```

LVs:	0	USED PPs:	0 (0 megabytes)
OPEN LVs:	0	QUORUM:	2
TOTAL PVs:	1	VG DESCRIPTORS:	2
STALE PVs:	0	STALE PPs	0
ACTIVE PVs:	1	AUTO ON:	yes
=====			

3. You can dynamically add disks to an existing volume group. For example, to add hdisk3 to vgsun02, proceed as follows:

a. Navigate to the volume group modification screen by highlighting and selecting the following items:

```
System Storage Management → Logical Volume Manager → Volume
Groups → Set Characteristics of a Volume Group
```

b. Select the Add a Physical Volume to a Volume Group function.

c. Press Esc-4 to list available volume group names. Highlight the volume group vgsun02.

d. Use the arrow keys to advance to the physical volume field. Press Esc-4 for a list of available physical volumes, arrow down to hdisk3, and press Enter to add this physical volume to the volume group.

```
=====
Add a Physical Volume to a Volume Group

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[Entry Fields]
* VOLUME GROUP name          [vgsun02] +
* PHYSICAL VOLUME names     [hdisk3] +
=====
```

4. Create a logical volume.

a. Navigate to the Add a Logical Volume screen by highlighting and selecting the following items:

```
System Storage Management → Logical Volume Manager → Logical
Volumes → Add a Logical Volume
```

- b. When you reach the Add a Logical Volume screen, press Esc-4 to display a list of available volume group names. Arrow down to highlight volume group vgsun02, and then press Enter.
- c. Type in the desired logical volume name (in this example, lvsun12).
- d. Arrow down to Number of LOGICAL PARTITIONS and supply the desired value. (This is equal to the desired space in Mbytes, divided by the size of physical partitions in MBytes.)
- e. Use the arrow keys to scroll to MAXIMUM NUMBER of LOGICAL PARTITIONS (on the next screen) and supply the desired number.
- f. Press the Enter key to create the logical volume.

```

=====
                          Add a Logical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                [Entry Fields]
Logical volume NAME                  [lvsun12]
VOLUME GROUP name                    vgsun02
Number of LOGICAL PARTITIONS         [300]                                #
PHYSICAL VOLUME names                []                                    +
Logical volume TYPE                  []                                    +
POSITION on physical volume          outer_middle                        +
RANGE of physical volumes            minimum                              +
MAXIMUM NUMBER of PHYSICAL VOLUMES  []                                    #
  to use for allocation
Number of COPIES of each logical     1                                    +
  partition
Mirror Write Consistency?            yes                                    +
Allocate each logical partition copy  yes                                    +
[MORE...10]
=====

```

5. Create a file system.

- a. Navigate to the Add a Journaled File System on a Previously Defined Logical Volume screen by highlighting and selecting the following items:

```

System Storage Management → File Systems → Add/Change/Show/Delete
File Systems → Journaled File Systems → Add a Journaled File
System on a Previously Defined Logical Volume

```

b. Press `Esc-4` to display a list of available logical volume names and select `lvsun12`.

c. Arrow down to `MOUNT POINT` and supply the desired path (here, `/A70002`).

Note – If the mount point does not exist, SMIT creates it.

d. Arrow down to `Mount AUTOMATICALLY at system restart?` and press the `tab` key until `yes` is displayed.

e. Create the file system by pressing `Enter`.

Note – This may take several minutes, depending on the size of the file system.

```
=====
Add Journaled File System on a Previously Defined Logical Volume

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

LOGICAL VOLUME name           [Entry Fields]
                               lvsun12           +
MOUNT POINT                   [/A70002]
Mount AUTOMATICALLY at system restart?  yes           +
PERMISSIONS                   read/write      +
Mount OPTIONS                 []              +
Start Disk Accounting?       no              +
Fragment Size (bytes)        4096           +
Number of bytes per inode    4096           +
Compression algorithm        no              +
=====
```

6. Mount the file system.

a. Select the following SMIT menu items:

```
System Storage Management → File Systems → Mount a File System
```

b. Use the arrow keys to highlight the new file system, and then press `Enter`.

c. To mount the file system, press Enter a second time.

```
=====
                          Mount a File System

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                     [Entry Fields]
FILE SYSTEM name                    [ /dev/A70002 ] +
DIRECTORY over which to mount      [ ] +
TYPE of file system                 +
FORCE the mount?                    no +
REMOTE NODE containing the file system [ ] +
to mount
Mount as a REMOVABLE file system?   no +
Mount as a READ-ONLY system?        no +
Disallow DEVICE access via this mount? no +
Disallow execution of SUID and sgid programs no +
in this file system?
=====
```

7. Exit from SMIT and check the file system.

```
IBM# fsck /A70002
```

A7000 Configurations and Limitations

This appendix provides three example A7000 configurations.

- Example 1: Maximum A7000 Configuration—page B-2
- Example 2: A7000 Configuration—page B-3
- Example 3: A7000 Configuration—page B-4

Example 1: Maximum A7000 Configuration

The A7000 can have a maximum of eight STE SCSI boards per side (node) with four ports per board equaling 32 STE ports. However, the system maximum external I/O connections per node is 16 ports. This system has:

- Two Nodes
- Six drawers for internal disks
- 32 external I/O connections (16 per node)

Available External I/O Connection Types

- BMC (installed)
- ESCON
- STE (installed)

324 Disk Drives (Internal)

- Disks are arranged in groups of six (six packs)
- 324 disks = 54 six packs; 27 six packs (162 disks) per node
- Disks are 9-Gbyte SCSI

14 Slots Per Node

- Two slots used for CPU
- One slot used for memory channel
- Seven slots for backend (internal) SCSI drives
- Three slots for front end STE, 4 ports/slot, total: 12 STE ports
- One slot for BMC, 4 ports/slot, total: 4 BMC ports

Example 2: A7000 Configuration

This system has:

- Two nodes
- Two drawers for internal disks

32 External I/O Connections (16 per Node)

- BMC
- ESCON (installed)
- STE (installed)

108 Internal Disk Drives

- Disks are arranged in groups of six (six packs)
- 108 disks = 18 six packs; 9 six packs (54 disks) per node
- Disks are a mix of 4-Gbyte and 9-Gbyte SCSI

14 Slots Per Node

- Two slots used for CPU
- One slot used for memory channel
- Three slots for backend (internal) SCSI drives
- Two slots for front end STE, 4 ports/slot, total: eight STE ports
- Four slots for ESCON, 2 ports/slot, total: eight ESCON ports
- Two empty slots (Not usable due to a maximum of 16 I/O ports per side; eight STE ports and eight ESCON ports)

Example 3: A7000 Configuration

This system has:

- Two nodes
- Four drawers for internal disks

32 external I/O connections (16 per node)

- BMC
- ESCON (installed)
- STE (installed)

240 Disk Drives (Internal)

- Disks are arranged in groups of six (six packs)
- 240 disks = 40 six packs; 20 six packs (120 disks) per node
- Disks are 4-Gbyte SCSI

14 Slots per node

- Two slots used for CPU
- One slot used for memory channel
- Five slots for backend (internal) SCSI drives
- Three slots for front end STE, 4 ports/slot, total: 12 STE ports
- Two slots for ESCON, 2 ports/slot, total: 4 ESCON ports
- One empty slot (Not usable due to a maximum of 16 I/O ports per side; 12 STE ports and four ESCON ports.)

Glossary

access functions	DataShare functions that interpret the non-native volume structures for a client system.
BLKL	DataShare keyword for the mainframe BLKSIZE.
BLKSIZE	The maximum length, in bytes, of a data block.
BMC	Block Mux Channel.
CKD	Count-Key-Data. A disk data recording technique used by mainframes.
clients	Systems that receive services from a server.
conversion functions	DataShare functions that modify the information format or content before presentation to another host.
DASDMGR	Direct Access Storage Device Manager, which are utilities used by Sun Enterprise Services to configure storage on the StorEdge A7000 system.
data replication	Reproduction of data on various platforms.
datasets	Mainframe terminology for files.
DataShare capabilities	Functions that are available through the DataShare Facility; that is, mainframe and open systems DataShare capabilities.
DSF	DataShare Facility.
ESCON	Enterprise System Connect, which is a high-speed connection from a mainframe to an A7000 system.
esoteric	Identifier by which a group of devices (usually tape and disk) are known. For example, you can use an esoteric name to request space from a specific pool of devices without having to know the volume serial number (VOLSER) of any specific device in the pool. Administrators can define security permissions at the esoteric level.
files	Open system terminology for mainframe datasets.

fixed block	A dataset structure in which blocks are a fixed length not exceeding 32K.
fixed record length	Data record containing fixed-length fields.
HDA	Head disk assembly.
IDRC	Improved Data Recording Capability, which is a hardware compression feature on 3480/3490 tape subsystems. For planning purposes, a compression ratio of 2 to 1 is usually assumed.
IEBGENER	IBM standard utility that copies files from one location to another. This utility is similar to the UNIX <code>dd</code> , <code>cat</code> , and <code>cp</code> commands.
ISO-9660	The standard on which CD-ROMs are based.
Julian date	Date format that has the year in positions one and two, and the day in positions 3 through 5. The day is 1 – 366 and is right-adjusted. Zeros are in the unused high-order positions.
logical disks	Virtual disks.
LRECL	Logical record length.
MVS	Multiple Virtual Storage (mainframe operating system).
OEMI	Other Equipment Manufacturer Interface.
open systems	Hardware and software that permit easy access to various vendor products providing application portability, scalability, and interoperability using approved standards. An open system is the opposite of a proprietary or vendor-specific implementation.
phantom header/tail	Certain open systems initiators require space at the beginning and/or end of each disk for header information; the required size and location (at the beginning or end of the disk) are system-specific. To prevent these initiators from writing their header information over the contents of a shared device, Open Backup prepends a “phantom” header and/or appends a “phantom” tail to the initiator’s view of the contents of the disk and store the initiator’s header information there. This area is always read/write and is independent of the access rights of the data shared region.
presentation functions	DataShare functions that enable delivery of accessed information to a specific client native format.
QSAM	Queued Sequential Access Method. A mainframe access method for flat files.
quiescent	No reads or writes being performed and all cached data destaged to disk.
record format	A mainframe term describing the physical format of a dataset: Fixed, fixed block, variable block, and variable block spanned.

SCSI Target Emulation	Hardware and software that allow an A7000 system to appear to open system initiators as one or more SCSI target disks.
servers	Systems that provide shared services to client systems.
SIMCKD	Simulation of Count-Key-Data.
SMS	System Managed Storage.
SNA	Systems Network Architecture
STE	SCSI Target Emulation.
subsystem	Combination of storage hardware (channel adapters, device adapters, CPUs, cache memory) and software that is independently powered and configured in an A7000. Typical A7000 storage platforms comprise two adjacent subsystems working as a pair to service a large array of disks. Subsystems are sometimes referred to as nodes.
virtual disk	The view provided by STE of the SCSI target disk from a SCSI initiator's perspective.
VM	Virtual Machine (mainframe operating system).
VOLSER	Volume serial number. A six character volume label used by the mainframe operating system.
volume	Mainframe terminology for a logical disk. From a mainframe perspective, it is a logical representation of a physical address.
VSE	Virtual Storage Extended (mainframe operating system).
VTOC	Volume Table of Contents. A mainframe equivalent of an inode or file allocation table.

Index

A

A7000

- architecture, 1-6
- client volume, 3-9
- configurations, B-1
- limitations, B-1

access functions, Glossary-1

applications

- examples, 5-1

architecture

- DataShare, 1-7
- illustration, 1-6
- StorEdge A7000, 1-6

assumptions, 3-2

B

backing up

- HP-UX, 4-16
- IBM AIX, 4-20
- Sun, 4-9
- Windows NT, 4-12

backup

- performance, optimizing, 2-15
- server
 - mainframe, 1-3, 2-1
 - Sun system, 1-4, 3-1
- window, 2-3, 3-2

bandwidth, 3-4

best practices, 5-1

- mainframe as backup server, 5-2

- Sun system as backup server, 5-7
- business continuation planning, 2-22

C

capacity

- data, 2-7
- planning, 3-4

client issues, 4-1

- device types, 4-3, 4-6
- disk hardware configurations, 4-6
- disk naming conventions, 4-7
- disk partitions, 4-3
- HP-UX, 4-15
- IBM AIX, 4-19
- Sun, 4-7
- Windows NT, 4-11

configuration

- A7000, B-1
- changes, 2-11
- disk hardware, 4-6
- updating (MVS I/O), 2-13
- updating (Open Backup), 2-13

configured disks

- Windows NT, A-7

conversion functions, Glossary-1

coordinating processing with the open system, 2-21

creating HP-UX file system, A-11

- D**
- DASDMGR, Glossary-1
 - data capacity, 2-7
 - Data Facility Data Set Services (DFDSS), 1-3
 - data protection, 2-10
 - DataShare
 - architecture, 1-7
 - architecture illustration, 1-7
 - determining the bandwidth, 3-4
 - device types
 - client issues, 4-3, 4-6
 - DFSMSDss JCL example, 2-18
 - disk
 - administrator
 - Windows NT, A-7, A-8
 - hardware configuration, 4-6
 - naming conventions
 - HP-UX, 4-15
 - IBM AIX, 4-19
 - Sun, 4-7
 - Windows NT, 4-11
 - partitions
 - client issues, 4-3
 - verifying Sun, A-3
- E**
- esoteric, 2-11, Glossary-1
 - examples
 - applications, 5-1
 - JCL, 2-16
- F**
- Fast Dump Restore (FDR), 1-3
 - FDR JCL example, 2-16
 - file systems
 - creating HP-UX, A-11
 - creating Windows NT, A-8
 - mounting Sun, A-3
 - UNIX, A-4, A-5
 - unmounting, A-5
 - fixed block, Glossary-2
 - formatting
 - HP-UX disk from command line, A-13
 - HP-UX disk with System Administration Manager (SAM), A-11
 - HP-UX disks, A-11
 - Sun disks, A-2
 - Windows NT disks, A-8
 - Windows NT partitions, A-8
- frequency
- planning, 3-5
- functions
- access, Glossary-1
 - conversion, Glossary-1
 - presentation, Glossary-2
- H**
- hardware, A-10
 - configurations, client issues, 4-6
 - scanning or rescanning Sun, A-2
 - HP-UX
 - backing up, 4-16
 - client issues, 4-15
 - creating a file system, A-11
 - disk naming conventions, 4-15
 - file system creation, A-11
 - formatting disks, A-11
 - mounting a file system, 4-16
 - partitioning disks, A-11
 - platform-specific information, A-10
 - restoring, 4-17
 - scan or rescan hardware, A-10
 - unmounting a file system, 4-16
- I**
- IBM AIX
 - backing up, 4-20
 - client issues, 4-19
 - creating the file system, A-19
 - disk naming conventions, 4-19
 - mounting a file system, 4-19
 - optimizing device attributes, A-19
 - restoring, 4-21
 - scanning or rescanning the hardware, A-18
 - unmounting a file system, 4-20
 - IEBGENER, Glossary-2
 - JCL example, 2-20

- implementation on a mainframe, 2-13
 - initializing, 2-14
 - restarting the system, 2-14
 - updating MVS I/O configuration, 2-13
 - updating Open Backup configuration, 2-13
 - varying mainframe volumes offline, 2-13
- ISO-9660, Glossary-2

J

- JCL examples, 2-16
 - DFSMSdss, 2-18
 - FDR, 2-16
 - IEBGENER, 2-20
- Julian date, Glossary-2

L

- limitations
 - A7000, B-1

M

- mainframe
 - assumptions, 2-4
 - backup server, 2-1
 - best practices, 5-2
 - implementation, 2-13
 - planning
 - data capacity, 2-7
 - data protection, 2-10
 - determining number of disks, 2-5
 - determining size of disks, 2-5
 - how to connect the open systems host volumes, 2-8
 - naming convention for the open systems host, 2-9
 - required esoteric device name changes, 2-11
 - schedule time with Sun Enterprise Services, 2-11
 - type and number of mainframe volumes, 2-6
 - worksheet, 2-23
- mounting a file system
 - HP-UX, 4-16
 - IBM AIX, 4-19
 - Sun, 4-8, A-3
 - Windows NT, 4-11

- MVS mainframe
 - backup server, 1-3

O

- Open Backup, using, 2-14, 3-10
 - coordinating processing with the open system, 2-21
 - JCL examples, 2-16
 - optimizing backup and restore performance, 2-15
 - quiescing the client system, 2-15
- optimizing backup and restore performance, 2-15

P

- partitioning
 - HP-UX disk from command line, A-13
 - HP-UX disk with System Administration Manager (SAM), A-11
 - HP-UX disks, A-11
 - Sun disks, A-2
- partitions
 - formatting, A-8
- performance
 - planning, 3-5
- phantom
 - header, 3-7, Glossary-2
 - tail, 3-7, Glossary-2
- planning
 - business continuation, 2-22
 - process, mainframe as backup server, 2-5
 - worksheet
 - mainframe backup server, 2-23
 - Sun backup server, 3-12
- platform-specific information, A-1
 - HP-UX, A-10
 - Windows NT, A-6
- presentation functions, Glossary-2
- primary partition
 - Windows NT, A-8
- product description, 1-1

- Q**
- quiescence, 4-2
 - quiescing the client system, 2-15
- R**
- record format, Glossary-2
 - recovery
 - planning, 3-5
 - restarting the system, 2-14
 - restore performance
 - optimizing, 2-15
 - restoring
 - elsewhere, 4-6
 - HP-UX, 4-17
 - IBM AIX, 4-21
 - Sun, 4-10
 - Windows NT, 4-13
- S**
- scan or rescan HP-UX, A-10
 - SCSI IDs
 - Windows NT, A-6
 - SCSI Target Emulation, Glossary-3
 - and DataShare Open Backup, 1-6
 - SIMCKD, Glossary-3
 - and DataShare Open Backup, 1-6
 - special considerations
 - data protection, 2-10
 - starting volumes
 - Sun, A-3
 - STE, Glossary-3
 - stopping volumes
 - Sun, A-5
 - StorEdge A7000
 - architecture, 1-6
 - configurations, B-1
 - limitations, B-1
 - subsystem, Glossary-3
 - Sun
 - backing up, 4-9
 - backup server, 1-4, 3-1
 - best practices, 5-7
 - phantom header, 3-7
 - phantom tail, 3-7
 - setup, 3-7
 - client issues, 4-7
 - disk naming conventions, 4-7
 - formatting disks, A-2
 - mounting a file system, 4-8
 - mounting file systems, A-3
 - partitioning disks, A-2
 - planning
 - capacity, 3-4
 - frequency, 3-5
 - performance, 3-5
 - recovery, 3-5
 - worksheet, 3-12
 - restoring, 4-10
 - scanning or rescanning hardware, A-2
 - Solaris operating environment, 4-7, A-2
 - starting volumes, A-3
 - stopping volumes, A-5
 - UNIX file system, A-3, A-4, A-5
 - unmounting a file system, 4-8
 - unmounting file systems, A-5
 - volume manager, A-4, A-5
 - System Administration Manager (SAM), A-11
- U**
- UFS, see UNIX file system, A-3
 - UNIX file system
 - Sun, A-3, A-4, A-5
 - unmounting a file system
 - HP-UX, 4-16
 - IBM AIX, 4-20
 - Sun, 4-8, A-5
 - Windows NT, 4-11
 - updating configuration, 2-13
 - using Open Backup, 2-14, 3-10
- V**
- varying mainframe volumes offline, 2-13
 - verify disk partitioning
 - Sun, A-3
 - virtual disk, Glossary-3
 - VOLSER, Glossary-3
 - volume manager

- Sun, A-4, A-5
- volumes
 - starting, A-3
 - stopping, A-5
- VTOC, Glossary-3

W

- Windows NT, A-8
 - backing up, 4-12
 - client issues, 4-11
 - configured disk, A-7
 - disk administrator, A-7, A-8
 - disk naming conventions, 4-11
 - formatting disks, A-8
 - formatting partitions, A-8
 - mounting a file system, 4-11
 - platform-specific information, A-6
 - primary partition, A-8
 - restoring, 4-13
 - SCSI IDs, A-6
 - unmounting a file system, 4-11

Y

- Y2K compliance, 1-8
- year 2000 compliance, 1-8

