

STORE and TurboSTORE/iX Products Manual

HP 3000 MPE/iX Computer Systems

Edition 2



**Manufacturing Part Number: B5151-90003
E0300**

U.S.A. March 2000

Notice

The information contained in this document is subject to change without notice.

Hewlett-Packard makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for direct, indirect, special, incidental or consequential damages in connection with the furnishing or use of this material.

Hewlett-Packard assumes no responsibility for the use or reliability of its software on equipment that is not furnished by Hewlett-Packard.

This document contains proprietary information which is protected by copyright. All rights reserved. Reproduction, adaptation, or translation without prior written permission is prohibited, except as allowed under the copyright laws.

Restricted Rights Legend

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c) (1) (ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013. Rights for non-DOD U.S. Government Departments and Agencies are as set forth in FAR 52.227-19 (c) (1,2).

Acknowledgments

UNIX is a registered trademark of The Open Group.

Hewlett-Packard Company
3000 Hanover Street
Palo Alto, CA 94304 U.S.A.

© Copyright 1996, 2000 by Hewlett-Packard Company

Contents

1. Introduction

System Backups	16
Product Structure and Feature Comparisons	17
Hardware Requirements	18
Backup Media	20
Backup Devices	20
Disaster Recovery Plans	21
Related Documentation	21

2. Managing Backups and Backup Media

Why Perform Backups?	23
Determining Your Backup Strategy	24
Using STORE and RESTORE	26
Helping Users Store and Restore Files	26
What Should You Backup?	29
Scheduling Backups	30
Full and Partial Backups	30
Selecting a Time for Backups	31
Using Unlabeled vs Labeled Media	31
Maintaining a Backup Library	32
Maintaining a Large Stock Of Backup Media	32
Keeping Backup Records	32
Storage of Backup Media	32

3. Preparing Storage Devices

Preparation Process	35
Selecting the Backup Method	37
Selecting Backup Device Order	37
Assigning Device File Names	38
Preparing Tapes and DDSs	40
Preparing Tape and DDS Devices	46
Single Tape Devices	46
Multiple Tape Devices	47
Using \$NULL as a STORE Device	52
Using Remote Devices for Backups	52
Storing to Disk Files	54

4. Preparing the System

Notifying Users of Pending Backup	58
Preventing New System Activity	59
Sending Warning Messages	60
Replying to Console Requests	61

Contents

Temporarily Suspending Jobs	62
Sending Final Warning Messages	63
Aborting Any Remaining Sessions	64
5. Storing Files	
Specifying Files in a STORE Command	66
Filesets	66
Specifying A Fileset	67
Using Wildcard Characters	68
MPE and HFS Naming Equivalences	69
Using Indirect Files	70
Using Filesets With the RENAME Option	72
Storing Symbolic Links and Device Links	72
Storing Database Files	73
Storing to Labeled Tapes	75
Using Previously Labeled Volumes	75
Using Unlabeled Volumes	75
Monitoring STORE Command Progress	77
Responding to Tape Requests	77
Replying to Tape Requests for a Multiple-Device Backup	77
Handling Media Errors	78
Responding to a NO WRITE RING Message	79
Progress Messages	79
Mounting Additional Tapes	79
6. STORE and TurboSTORE/iX Command Options	
STORE Options	83
Listing the Files Stored (SHOW)	83
Selecting an Error Recovery Method (ONERROR)	89
Specifying Files by Date (DATE)	90
Removing Stored Files from the System: Archiving Files (PURGE)	91
Displaying Progress Messages (PROGRESS)	92
System Directory Backup (DIRECTORY)	92
Storing Files in Compatible Formats (TRANSPORT)	93
Specifying the Number of Files to Store (FILES)	95
Managing File Security (COPYACD and NOACD)	96
Storing Files with Certain File Codes (FCRANGE)	96
Improving Storing Efficiency by Increasing Record Size (MAXTAPEBUF)	96
Notification of File Availability (NOTIFY)	97
Additional Backup Information (STATISTICS)	98
Storing Files from Volume Sets (ONVS)	103
Storing Files from Split Volumes (SPLITVS)	103

Contents

Renaming Files (RENAME)	104
Specifying HFS Files (TREE and NOTREE)	109
Storing Database Files Using PARTIALDB or FULLDB	109
Creating a Disk Directory File (STOREDIRECTORY and NOSTOREDIRECTORY) . . .	110
Storing Files in Interleave Format (INTER)	112
TurboSTORE/iX II Options	113
Compressing Data (COMPRESS)	113
TurboSTORE/iX 7x24 True-Online Backup Options	114
Storing Files Online (ONLINE)	114
Specifying Where to Place Log Files (LOGVOLSET)	117

7. TurboSTORE/iX 7x24 True-Online Backup

Using TurboSTORE/iX 7x24 True-Online Backup	120
The Sync Point	120
Sync Point at the Beginning of Backup	121
Sync Point at the End of Backup	121
Consistent Backups	123
Physical Consistency	123
Logical Consistency	123
How Shadow Logging Works	125
When the Sync Point is at the Beginning	125
When the Sync Point is at the End	126
Logging File Attribute Modifications	127
File Creations and Purges	127
Other File Attributes	127
TurboSTORE/iX 7x24 True-Online Backup Options	128
Creating a TurboSTORE/iX 7x24 True-Online Backup (ONLINE)	128
Initiating a TurboSTORE/iX 7x24 True-Online Backup (START and END)	128
Specifying Time	128
Prompting During Backup (ASK)	128
Roll-forward Recovery of TurboIMAGE/XL Databases (DBRECOV)	130
Storing Part of a Database (PARTIALDB)	131
Creating a Disk Directory File (STOREDIRECTORY and NOSTOREDIRECTORY) . . .	131
Restoring TurboSTORE/iX 7x24 True-Online Backups	132
Restoring with a Disk Directory File (STOREDIRECTORY)	132

8. Handling Media Requests

Preparing the Backup Device	136
Checking the Console	136
Replying to a Media Request	137
Pending Media Requests	138
Denying a Media Request	138

Contents

Using REPLY and RECALL	138
Automatically Assigning Devices	139
9. When STORE is Complete	
Checking File Listings for Errors	142
Verifying That STORE Files Are Error Free Using VSTORE	143
Using NODECOMPRESS	145
Retaining the STORE Tapes	145
Allowing Users Back on the System	146
10. Backing Up the System Configuration	
SYSGEN Backup Capabilities	150
Running SYSGEN	151
The System Backup Command	152
Creating a System Load Tape	153
Combined SLT and Store Set	155
11. Restoring Files	
Preparing to Restore Files	159
Locate the Necessary File or Files	159
Checking for Duplicate Filenames	159
Preparing the Media and the Backup Device	159
Referencing Files in a RESTORE Command	160
Restoring from Labeled Tapes	161
Restoring from Single Devices	161
Restoring from Multiple Devices (RESTORESET)	162
Sequential Devices	162
Parallel Devices	163
Restoring Files from a Disk File	164
Specifying the Disk File	164
Restoring From Multiple Disk Files	164
Restoring from Parallel Disk Files	165
Verifying Disk Backups	166
Issuing a RESTORE Command	167
Restoring Different Format Tapes	168
RESTORE Command Options	169
Listing Restored Files (SHOW)	169
Selecting a RESTORE Error Recovery Method (ONERROR)	170
Restoring Files to the Correct Group, Account, and Creator (CREATE)	171
Changing a File's GID (GID)	173
Overwriting or Retaining Disk Files (KEEP, NOKEEP)	173
Specifying Modification and Last Access Date (OLDDATE, NEWDATE)	173

Contents

Restoring Directory Information (DIRECTORY)	174
Finding Out What Is on Your Tape (LISTDIR)	174
Displaying Progress Messages (PROGRESS)	175
Restoring Files with Certain File Codes (FCRANGE)	176
Specifying a Maximum Number of Files to Restore (FILES)	176
Restoring Files to a Specific Device (DEV)	176
Restoring Files to Specific Volumes (VOL, VOLCLASS, VOLSET)	176
Modifying ACDs (COPYACD, NOACD)	177
Specifying HFS Files (TREE, NOTREE)	177
Restoring Part of a Database (PARTIALDB)	178
Restoring using a disk directory file (STOREDIRECTORY)	178
Monitoring RESTORE Command Progress	180
Responding to Tape Requests	180
Replying to Tape Requests for a Multiple-Device RESTORE	180
Mounting Additional Tapes	180
Using the PROGRESS Option	181
When RESTORE is Complete	182
12. Using Backup Media to Recover from a System Failure	
Reloading Your System Configuration	184
Restoring Files from Partial Backup Tapes	185
Restoring Files from Full Backup Tapes	186
13. Disaster Recovery	
Recovery Facilities	188
Private Recovery Facilities: Cold Sites and Hot Sites	188
Mutual Backup Arrangements	188
Hewlett-Packard Disaster Recovery Services	188
Commercial Backup Companies	189
Computer Service Bureaus	189
Planning for Disaster Recovery	190
The Recovery Facility	190
System Compatibility	190
Disaster Notification and Response Time	190
Storage of Backup Tapes	191
Testing Your Recovery System	192
A. Error Messages	
Determining the Cause of STORE, RESTORE and VSTORE Errors	194
Viewing the Value of STOREJCW	194
Determining the Cause of an Error	194
Using the STORE_ERROR Variable	196

B. Command Syntax

STORE Syntax198
RESTORE Syntax199
VSTORE Syntax200

C. Transferring Optical STORE Backups onto Standard Tapes

Figures

Figure 2-1.. Sample Backup	30
Figure 3-1.. Multiple-Device Planning Chart	36
Figure 3-2.. A Correctly Labeled Tape	42
Figure 3-3.. Preparing Tapes for a Single Device	42
Figure 3-4.. Preparing Tapes for Parallel Devices	43
Figure 3-5.. Preparing Tapes for Sequential Devices	44
Figure 3-6.. Preparing Tapes for Parallel Device Pools	45
Figure 3-7.. Single-Device Backup	46
Figure 3-8.. Sequential Backup	48
Figure 3-9.. Parallel Backup	50
Figure 3-10.. Parallel Device Pools	51
Figure 4-1.. Aborted Session System Response	64
Figure 6-1.. Short File Display (MPE Format)	84
Figure 6-2.. Short File Display (HFS Format)	84
Figure 6-3.. Short File Display with Wrapped HFS Names	84
Figure 6-4.. Long File Display (MPE Format)	85
Figure 6-5.. Long File Display (HFS Format)	86
Figure 6-6.. File Date Information (MPE Format)	86
Figure 6-7.. File Date Information (HFS Format)	87
Figure 6-8.. File Security Information	87
Figure 6-9.. Names Only Format	88
Figure 7-1.. Backup Created with Sync Point at the Beginning	125
Figure 7-2.. Backup Created with Sync Point at the End	126
Figure 8-1.. Job Session Description	137
Figure 9-1.. Listing Files Copied to Tape	142
Figure 9-2.. Listing Suspended Jobs	146
Figure 9-3.. Listing Resumed Jobs	147
Figure 10-1.. SYSGEN First-Level Command Listing	151
Figure 10-2.. Creating a System Load Tape Sample	153
Figure 11-1.. Sample RESTORE Command and Output	175
Figure 11-2.. Listing Restored Files	182

Tables

Table 1-1.. Product Structure Change	17
Table 1-2.. Feature Comparison	17
Table 2-1.. Task Roadmap	27
Table 6-1.. Maximum Tape Record Sizes	97
Table 6-2.. Required Capabilities for RENAME	104
Table 6-3.. Creator File Translations	108
Table A-1.. STOREJCW Values	195

Preface

This book, the *STORE and TurboSTORE/iX Products Manual*, describes how to use the Fundamental Operating System (FOS) STORE product and the following TurboSTORE products:

Product Number	Product Description
-----------------------	----------------------------

FOS	STORE
B5151AA	TurboSTORE/iX II
B5152AA	TurboSTORE/iX 7x24 True-Online Backup

NOTE The product structure of the STORE and TurboSTORE/iX products has changed for MPE/iX Release 5.5. Some of the features have been placed or kept in the STORE/iX program in the MPE/iX Fundamental Operating System (FOS), and the rest have been consolidated into the two TurboSTORE/iX products listed above.

See Table 1-1. in Chapter 1 to see how the product structure has changed, and in Table 1-2. Chapter 1 for a comparison of features in the new products.

MPE/iX, MultiProgramming Executive with Integrated POSIX, is the latest in a series of forward-compatible operating systems for the HP 3000 line of computers.

In HP documentation and in talking with HP 3000 users, you will encounter references to MPE XL, the direct predecessor of MPE/iX. MPE/iX is a superset of MPE XL. All programs written for MPE XL will run without change under MPE/iX. You can continue to use MPE XL system documentation, although it may not refer to features added to the operating system to support POSIX (for example, hierarchical directories).

Finally, you may encounter references to MPE V, which is the operating system for HP 3000s, not based on PA-RISC architecture. MPE V software can be run on the PA-RISC (Series 900) HP 3000s in what is known as *compatibility mode*.

This book, the *STORE and TurboSTORE/iX Products Manual* describes the processes and procedures for using the STORE, RESTORE, and VSTORE commands. Specific directions for using your hardware is not described in this book. Refer to your hardware documentation.

This book is divided into 13 chapters and 3 appendices as follows:

- Chapter 1** **Introduction**, provides basic backup concepts and lists hardware requirements.
- Chapter 2** **Managing Backups and Backup Media**, describes planning, scheduling, and maintaining backups.
- Chapter 3** **Preparing Storage Devices**, provides procedures for using backup devices.

- Chapter 4** **Preparing the System**, provides procedures for notifying the users of the backup process.
- Chapter 5** **Storing Files**, provides procedures for using the `STORE` command to store files and backup media.
- Chapter 6** **STORE and TurboSTORE Options**, describes each of the `STORE` and `TurboSTORE` options.
- Chapter 7** **TurboSTORE/iX 7x24 True-Online Backup**, describes the operation and options of the new `TurboSTORE/iX 7x24 True-Online Backup` product.
- Chapter 8** **Handling Media Requests**, describes how to use backup media.
- Chapter 9** **When STORE is Complete**, provides procedures for verifying that the `STORE` process is complete.
- Chapter 10** **Backing Up the System Configuration**, describes how to use the `STORE` command and the `SYSGEN` utility to perform a complete system backup.
- Chapter 11** **Restoring Files**, describes how to use the `RESTORE` command to restore files to your system. The chapter includes descriptions of the `RESTORE` options.
- Chapter 12** **Using Backup Media to Recover from a System Failure**, describes how to restore your system from backup media.
- Chapter 13** **Disaster Recovery**, describes strategies for protecting your computer operations from a variety of disasters.
- Appendix A** **Error Messages**, describes how to use the `STOREJCW` command to identify the cause of an error.
- Appendix B** **Command Syntax**, lists the `STORE`, `RESTORE`, and `VSTORE` syntax.
- Appendix C** **Transferring Optical STORE Backups onto Standard Tapes**, describes how to use the `STORCOPY` utility for transferring optical `STORE` backups onto standard tape.

1 Introduction

This manual presents important information for system managers and operators on all aspects of storing files for backup, archive, or transfer. It describes how to use the `STORE` command and the TurboSTORE/iX products to save files and restore them back to the system. It also describes important methods for safeguarding the information in your system. More specifically, the manual describes how to do the following:

- Maintain your backup media and manage backup devices.
- Plan and manage system backups.
- Prepare the system for a backup.
- Perform a file and system directory backup with the `STORE` program.
- Back up the system configuration with the `SYSGEN` program.
- Use backup media to restore files, and to recover from data loss, file corruption, or a major system failure, with the `RESTORE` program and/or the `ISL INSTALL` utility.
- Prepare and implement a disaster recovery plan.

The remainder of this chapter defines important terms, summarizes information on protecting files from unexpected problems, and outlines the basic steps for storing and restoring files. The sections in the chapter are:

- System Backups
- Product Structure and Feature Comparisons
- Hardware Requirements
- Backup Media
- Backup Devices
- Disaster Recovery Plans
- Related Documentation

System Backups

Storing all of the information on your system to backup media is called backing up your system. You back up your system so that you have a copy of the information stored on it should it experience a major failure that causes some or all of your files to be lost. If your system experiences a failure, you can reload that information from your backup media. Refer to subsequent chapters for complete instructions on performing system backups and restoring files from backup media after a system failure.

Planning, managing, and performing necessary system backups are, perhaps, a system administrator's/operator's most important tasks. Regularly scheduled, properly performed system backups guarantee that, in the event of an unexpected system mishap, all information present on the system at the time of the last backup is currently available.

If a system failure does occur, the amount of data lost depends directly upon three factors:

- How recently the last backup was performed.
- The level of system activity since the last system backup.
- The extent of the damage.

As a system administrator, you are responsible for establishing a backup schedule that best meets the needs of your organization and all system users. For example, the schedule may include saving all files in a particular account every day, or more commonly, copying only those files that have been added or modified since the last complete system backup.

You or a member of your staff are probably also responsible for actually performing the backups. In addition to user files, you must be sure to back up all system files since you may need to restore these files from backup media if the system experiences problems.

Product Structure and Feature Comparisons

The following table, Table 1-1, shows how the previous TurboSTORE/iX product structure for MPE/iX Release 5.0 has been incorporated into the current product structure for MPE/iX Release 6.5.

Table 1-1. Product Structure Change

5.0 Product Structure			6.5 Product Structure	
Product #	Product Description		Product #	Product Description
30319A	TurboSTORE/iX I	====>	N/A	Part of FOS
36387A	TurboSTORE/iX II	====>	B5151AA	TurboSTORE/iX II
36397A	TurboSTORE/iX II with Support for Optical Disk	====>	B5151AA	TurboSTORE/iX II
36388A	TurboSTORE/iX II with Online Backup	====>	B5152AA	TurboSTORE/iX 7x24 True-Online Backup
36398A	TurboSTORE/iX II with Support for Online Backup and Optical Disk	====>	B5152AA	TurboSTORE/iX 7x24 True-Online Backup

NOTE This manual, *STORE and TurboSTORE/iX Products Manual*, and product structure apply only to STORE and TurboSTORE programs running on the MPE/iX 6.5 and later operating systems. If you are running an older version of MPE/iX, refer to the previous manual, *STORE and TurboSTORE/iX Manual*.

These products contain features that can help decrease the amount of time needed for backups. They also contain support for other backup devices, such as normal disks.

With all versions of STORE/iX and TurboSTORE/iX, you use the STORE and RESTORE commands to back up files to media and to restore files to disk. Table 1-1, compares features of STORE/iX, TurboSTORE/iX II, and TurboSTORE/iX 7x24 True-Online Backup.

Table 1-2. Feature Comparison

Feature	STORE/ iX FOS	TurboSTORE/iX II B5151AA	TurboSTORE/iX 7x24 True-Online Backup B5152AA
Multiple Store Devices (STORESET)	YES	YES	YES
File Interleaving (INTER)	YES	YES	YES
Data Compression (COMPRESS)	NO	YES	YES

Table 1-2. Feature Comparison

Feature	STORE/ iX FOS	TurboSTORE/iX II B5151AA	TurboSTORE/iX 7x24 True-Online Backup B5152AA
Parallel Restore (RESTORESET)	NO	YES	YES
Optical Device (MOSET)	NO	NO	NO
Online Backup (ONLINE)	NO	NO	YES
7x24 True-Online Backup (ONLINE=START, END)	NO	NO	YES
Store to Disk	NO	NO	YES
Restore from Disk	NO	NO	YES
Labeled Tapes	YES	YES	YES

NOTE The TurboSTORE/iX 7x24 True-Online Backup product **does not** imply 24x7 (24 hours, 7 days a week) support. 7x24 True-Online Backup represents the availability of the data that this product offers by not having to shut down business critical applications for backups. Customers should consult their SE for levels of support with this product.

Hardware Requirements

All versions of TurboSTORE/iX work with any hardware configuration supported by the MPE/iX system.

When using the basic Store and Restore functionality available in the fundamental operating system (FOS), the minimum hardware requirement is at least one disk drive and at least one tape drive. The tape drive can be a magnetic tape drive, DDS drive, or StorageTek cartridge tape drive.

FOS STORE also allows you to store to more than one device at a time, either serially or in parallel. When using parallel or serial devices for a backup, up to 16 devices can be used in parallel. Within each parallel device pool, up to 8 serial devices are supported. Each parallel device will run as a separate process. Multiple processor systems, therefore, will see an even greater increase in performance when using parallel backups. A general rule is to use one less parallel device than there are processors. For a four-way MP system, this would mean three parallel devices.

NOTE From MPE/iX Release 6.5 onwards, TurboSTORE/iX II and TurboSTORE/iX 7x24 True-Online Backup do not support magneto-optical devices..

When using TurboSTORE/iX 7x24 True-Online Backup, you can back up to and restore from disk files. For the best performance, it is recommended that the store-to-disk file(s) be located on a different device than the files being stored or restored. However, this is not required. Storing to more than one disk file in parallel is supported. Again, for the best performance, each parallel disk file should reside on a different disk or different volume set. Serial device "pools" are not supported for disk files.

Although Store allows you to use a variety of backup devices, you cannot mix backup devices within a single backup. If you want to use multiple devices for a backup, they must all be one type. SCSI and HPIB DDS devices and SCSI and HPIB magnetic tape devices can be used together, however. Disk files cannot be mixed with any tape devices.

For more details on configuring your system for maximum backup performance, refer to the *HP3000 900 Series Computer System Configuration Guide*. Your optimal hardware configuration will depend on your backup requirements and type of system.

Backup Media

Backup media can include reel-to-reel magnetic tapes, DDS tapes, 3480 cartridges, and disk files. Backup media provides an economical way to maintain duplicates of important files on your system.

The process of copying a file to backup media is called storing the file. The process of copying a file from backup media is called restoring the file.

As a system administrator/operator, you frequently STORE and RESTORE files. You might, for example:

- Store all files on the system (backup) in case one or more files become corrupted, either accidentally, intentionally, or by a malfunction.
- Store an important file in order to create a duplicate for a user.
- Permanently store several files offline by storing them to media and deleting them from disk (archive).
- Transfer files from one system to another by storing them from one system and restoring them on to the other.

Backup Devices

Backup devices are system peripherals that allow you to write information to, and read information from, backup media. The backup devices currently supported by MPE/iX are:

- 1/2 inch tape drives
- DDS drives
- 3480 compatible drives
- Disk files

Because only one person can use a backup device at a time, it is your responsibility to monitor backup devices, grant users access to them, and help them load and unload backup media. Chapter 3, "Preparing Storage Devices," describes how to manage backup devices.

Disaster Recovery Plans

Regular system backups protect the information in your system from irretrievable loss by providing a copy of the information. However, natural and man-made problems, such as smoke, noxious fumes, explosion, flood, earthquake, and sabotage, can still threaten both your backup media and your system hardware. A disaster recovery plan protects your organization from such potential disablement by establishing procedures for duplicating and safely storing backup media and maintaining a site for "substitute" computer hardware. Although implementing a disaster recovery plan is time-consuming and costly, saving your system from an unexpected disaster is almost always worth the effort. Refer to Chapter 13, "Disaster Recovery," for more detail.

Related Documentation

Before you begin to plan and perform system backups, you should have a working knowledge of several system administrator functions. In particular, performing a system backup draws from information presented in the following manuals:

- *System Startup, Configuration, and Shutdown Reference Manual*
- *MPE/iX Commands Reference Manual*
- *Performing System Management Tasks*
- *Performing System Operation Tasks*
- *MPE/iX System Utilities Reference Manual*

Other documentation that provides additional information related to various aspects of the store and restore processes include:

- *HP3000 900 Series Computer System Configuration Guide*

The information presented in this book, *STORE and TurboSTORE/iX Products Manual*, does not describe how to use any specific hardware to perform the store and restore procedures. Refer to your hardware documentation for that specific information.

2 Managing Backups and Backup Media

This chapter describes the purpose and process of performing regular backups of your files and systems. The sections in this chapter are:

- Why Perform Backups?
- TurboSTORE Backup Features
- STORE and RESTORE Process
- What Should You Backup?
- Scheduling Backups
- Preparing to STORE Files
- Using Unlabeled vs Labeled Media
- Maintaining a Backup Library

Why Perform Backups?

Selecting backup methods and scheduling regular backups are important aspects of a system administrator's job. Standard procedures reduce the time necessary to perform backups and guarantee that the most recent system information is always available. A regular backup schedule reduces system interruption.

An important part of planning and performing system backups is making sure that the resources you need are available and well managed. Maintain a supply of backup media and labels and establish procedures for identifying and retaining backup media. Examples of these procedures are explained in the pages that follow.

Determining Your Backup Strategy

The TurboSTORE programs allow you to store files to a single device or to multiple devices.

The best way to use your backup devices depends upon the characteristics of the files that you are storing, your needs for compatibility with other versions of MPE, and your system configuration.

The characteristics of the set of files help determine choices (or options) for storing them. Use the following guidelines to select the most appropriate method:

- How many files do you need to back up?

Generally, the more files in the fileset, it is to your advantage to use multiple devices or a high capacity storage device such as a magneto-optical disk device.

- If you are storing only a few files of moderate size, one backup device such as a tape drive or DDS-format device is sufficient.
- If you are storing many files (for example, all the files in an account, or a smaller number of very large files), a multiple-device backup is best.
- If you need to do small, fast application backups, you may want to use the store-to-disk feature of TurboSTORE/iX 7x24 True-Online Backup (B5152AA). This allows you to quickly backup key applications to disk files.
- Do the backup files need to be compatible with another system?

Compatibility with other versions of MPE limits your choices. If you need to create STORE tapes that are compatible with MPE V/E, you must use the `TRANSPORT` option. This restricts you to a subset of the available STORE commands, and allows you to only store to one device. If you need to take files from a MPE/iX 5.0 or later system to a pre-5.0 system, you should use the `TRANSPORT=MPEXL` option. This ensures that various POSIX features are correctly translated to a pre-POSIX version of MPE.

- Does your system configuration include multiple backup devices?

If you do not need to create backups compatible with older versions of MPE, then your system configuration is the most important determining factor of which multiple-device method to use. Use the following guidelines to select a method:

- Which and how many backup devices do you have?

If you have two or three like backup devices, use either parallel or sequential devices. If you have at least four devices, use two or more parallel device pools.

- Do you intend to use labeled tapes?

Labeled tapes can be used with a single device and parallel devices, but not with sequential devices. Unlabeled tapes can be used with single, parallel, and sequential devices.

- Do you have user volumesets or mirrored volume sets on your system?

If you use volume sets for application and volume management, then you may want to use the `ONVS` option of `STORE`. This option allows you to back up all files on a particular volumeset, including the system directory information for the volumeset. Doing backups on a volumeset basis can help reduce downtime by breaking up a large backup into individual volumeset backups. Recovery time is also increased, since you can restore from backups that only contain files for a particular volumeset.

If you are using mirrored disks, and you do not have the TurboSTORE/iX 7x24 True-Online Backup product, you can use the `SPLITVTS` option of `STORE` to back up a split copy of your mirrored set.

However, if you are using TurboSTORE/iX 7x24 True-Online Backup, the recommended way of performing mirrored disk backups is to use either the `ONLINE=START` or `ONLINE=END` option. TurboSTORE/iX 7x24 True-Online Backup allows you to back up your volumeset without losing the protection of the mirrored disks, and without closing down files.

- Do you have critical applications that need to be available 24 hours a day?

If your system environment includes applications that need to be available on a 24-hour basis, you should use TurboSTORE/iX 7x24 True-Online Backup. This product provides three different online backup options:

- `ONLINE`, which requires files to be closed at the start of the backup.
- `ONLINE=START` or `ONLINE=END`, which allow files and applications to remain open throughout the entire backup.

Using any of the online backup options means that you do not have to dedicate the system to doing backups. Using the new 7x24 True-Online backup option allows you to backup files and applications while users are accessing them. Using the normal online option requires files being backed up to be closed for a short time at the start of the backup. Normal access can be resumed once the actual backup begins.

The TurboSTORE/iX 7x24 True-Online Backup options allow you to schedule the *sync point* for the backup at either the beginning of the backup (`ONLINE=START`) or at the end of the backup (`ONLINE=END`). The backup's sync point is the point when the state of all files being backed up is synchronized. When the backup is restored, all files are restored to the state they were in at the time of the sync point. TurboSTORE also synchronizes any TurboIMAGE or ALLBASE/SQL databases being stored, ensuring that consistent copies of the databases are written to the backup.

See Chapter 7 for more information on using the new TurboSTORE/iX 7x24 True-Online Backup feature.

Using STORE and RESTORE

The backup (STORE) and recovery (RESTORE) processes have similar steps. Only the specific commands and command options differ. The structure of your organization determines who performs the STORE and RESTORE steps. The basic steps include the following:

1. Determine what you want to store or restore. Refer to "What Should You Backup?" in this chapter.
2. Identify the storage device. Refer to Chapter 3, "Preparing Storage Devices."
3. Mount the media on the device. Refer to Chapter 3, "Preparing Storage Devices."
4. Prepare the system for store processes. Refer to Chapter 4, "Preparing the System." for preparing to STORE. Refer to Chapter 11, "Restoring Files." for preparing to RESTORE.
5. Issue the STORE or RESTORE command with the appropriate options. Refer to Chapter 5, "Storing Files" and Chapter 6, "STORE and TurboSTORE Command Options" or Chapter 11, "Restoring Files."
6. Reply to the console (media allocation) request. This may not apply if you are using labeled media. Refer to Chapter 8, "Handling Media Requests."
7. Remove the storage media. Refer to Chapter 8, "When STORE is Complete," and Chapter 11, "Restoring Files."

Helping Users Store and Restore Files

Storing and restoring files is not limited to system personnel. Account managers can store and restore any file in their accounts (except files with negative file codes) and all users can store and restore any file in their logon groups and accounts. While these other users can enter commands to store and restore at their terminals, they usually do not have access to backup devices.

The steps described below summarize how the operator and a user work together to store and restore files. The order in which you perform steps of the procedure may vary. For example, you may see a media request before the user hands you the backup media, instead of the other way around.

1. The user gives you empty backup media onto which to store disk files, or the user gives you the STORE media from which to restore files to disk.
2. You mount the media and prepare the device.
3. The user enters the FILE and STORE or RESTORE commands.
4. You check the console for a media request.
5. If the system does not automatically answer the media request, you must reply to it.
6. When the process is complete, remove the media and return it to the user.
7. As part of your regular routine of retrieving, separating, and distributing reports, collect the report describing the stored or restored files. Either file the report in the appropriate pickup bin, or keep it until the user comes to get it from you.

Table 2-1. provides a roadmap for system managers and operators of the chapters that pertain to their particular tasks.

Table 2-1. Task Roadmap

Chapter Description	Tasks			
	Define all backup and recovery procedures	Translate requirements into TurboSTORE commands	Create one or more job streams	Perform backup and recovery
1 "Introduction" - Basic backup concepts and hardware requirements.	X			
2 "Managing Backups and Backup Media" - Planning, scheduling, and maintaining backups.	X			
3 "Preparing Storage Devices" - Procedures for using backup devices.	X	X		
4 "Preparing the System" - Procedures for notifying users of backup process.		X		X
5 "Storing Files" - Procedures for using STORE to store files and backup media.			X	X
6 "STORE and TurboSTORE Command Options" - STORE and TurboSTORE option descriptions.		X	X	
7 "TurboSTORE/iX 7x24 True-Online Backup" - Describes the new 7x24 true-online backup product.	X	X	X	
8 "Handling Media Requests" - How to use backup media.				X

Table 2-1. Task Roadmap

Chapter Description	Tasks			
9 "When STORE is Complete" - Procedures for verifying STORE process is complete.	X		X	X
10 "Backing Up the System Configuration" - How to perform a complete system backup with STORE and SYSGEN utility.	X		X	X
11 "Restoring Files" - How to use RESTORE and its options.	X	X	X	X
12 "Using Backup Media to Recover from a System Failure" - How to restore a system from backup media.				X
13 "Disaster Recovery" - Strategies for protecting against system disasters.	X			

NOTE Non-Large files are those files with a file size less than four Giga Bytes, and Large files are files with a file size greater than or equal to four Giga Bytes and up to one Tera Byte.

NOTE From MPE/iX 6.5 onwards, TurboSTORE/iX supports store and restore of Large files. As of MPE/iX 6.5 release, the types of Large files supported are Fixed ASXII and KSAM 64. Please note that these files cannot be restored on a pre-6.5 system.

Forward Compatibility: Any media stored using Turbo Store from a previous version of MPE/iX can be restored using the TurboSTORE/iX 6.5.

Backward Compatibility: Any media stored using TurboSTORE/iX 6.5 can be restored onto a pre-6.5 system. However, if the store media containing Large Files is restored on a pre-6.5 system, only Non-Lagre files will be restored. A pre-6.5 restore skips the large files with the following error:

"<file>.<group>.<account> NOT RESTORED: MEDIA RECORD HEADER BLOCK CORRUPTED/INVALID"

What Should You Backup?

To completely back up your system, you should copy three types of information:

- User and system files. User files are usually the largest part of a system backup. They cannot otherwise be duplicated, except from hardcopy. Since, in total, they can be the product of several months or even years of labor, user files can be considered the most important element within your computer system. All system files are in groups within the `SYS` account; most system files are owned by the `PUB` group of the `SYS` account.
- The file system directory. The file system directory contains your account, group, and user structure, plus all associated capabilities. Each volume set has its own directory that must be saved to provide a complete backup.
- The current system configuration. The system configuration includes system tables, queues, device I/O attachments, and other parameters.

The various types of data are backed up using different mechanisms:

- System configuration information is backed up on a system load tape (SLT).
- User and system files, and the file system directory are backed up on `STORE` media, which can include tape, DDS, or disk.

Backing up system configuration files separately from other files is to your advantage because user, system, and system directory files usually change often while system configuration files change infrequently. Also, with the capability of backing up each type of file separately, you can avoid backing up any inadvertent corruption of system files while still backing up your other files.

Use the `SYSGEN` program to create a system load backup to supplement a `STORE` backup. The `SYSGEN` program enables you to back up the system configuration. You cannot back up operating system and configuration files with the `STORE` program. When you use `STORE` to perform system backups, you must still use the `SYSGEN` program to create a system load tape from which you can reload your system configuration.

To put SLT and store filesets on a single tape, you can use the combined SLT and Storeset functionality. Refer to Chapter 10, "Backing Up the System Configuration," for details of using the `SYSGEN` utility and combining an SLT and Store set.

Scheduling Backups

There are two important considerations when planning a backup schedule:

- Plan full and partial backups to efficiently duplicate the most recent system information.
- Schedule a regular time for performing backups.

Full and Partial Backups

There are two ways to perform a system backup. A full backup copies all files stored in your system whether or not any of the information in them has changed since the last backup. A partial, or incremental, backup copies only the files that have changed since the last full backup. This can be done with the `DATE` parameter of the `STORE` command. A well-planned backup schedule combines both methods. It usually calls for a full backup once a week and a partial backup each remaining workday. If you follow such a schedule, you would lose no more than 24 hours worth of information should the system experience problems.

Your requirements for recovery time should also be considered. Although it takes less time to make partial backups of your data between full backups, it also takes longer to restore them, should you need to. When restoring from partial backups, the last complete full backup must be restored first. Then, each subsequent partial backup must be restored. This is in contrast to spending more time each night on a full backup, but only needing to restore files once. You should examine your recovery time needs and expectations before deciding on a backup strategy.

Figure 2-1. shows a sample backup schedule that combines both full and partial system backups.

Figure 2-1. Sample Backup

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
					Full Backup	
	Partial Backup	Partial Backup	Partial Backup	Partial Backup	Full Backup	
	Partial Backup	Partial Backup	Partial Backup	Partial Backup	Full Backup	
	Partial Backup	Partial Backup	Partial Backup	Partial Backup	Full Backup	
	Partial Backup	Partial Backup	Partial Backup	Partial Backup	Full Backup	

LG200021_046a

Selecting a Time for Backups

It is best to perform backups at a time when people are not normally using the system. Often, system administrators choose to perform backups at the close of the workday or later in the evening. It is also a good idea to schedule the backup for the same time every day so users know when to expect it and can plan ahead.

Be sure to schedule enough time for each backup. Full backups take more time and more media than partial backups. Performing a partial backup a day or so after a full backup takes less time and less backup media than a partial backup performed almost a week later.

As the week progresses, users will have modified or added more files since the last full backup, and more new information must be duplicated. If you perform a full backup on Friday, the first partial backup on Monday may take only one tape, for example, and not much time. By Thursday, however, the day before the next scheduled full backup, a partial backup requires several tapes, for example, and considerably more time, but still less than a full backup.

The TurboSTORE/iX 7x24 True-Online Backup `ONLINE` option allows you to back up critical data and applications that need to be available at all times. Using just the `ONLINE` option allows users to access files *after* a short period of inaccessibility at the beginning of the backup. However, using the new `ONLINE=START` or `ONLINE=END` options, users can continue to access files, uninterrupted, throughout the backup.

Refer to Chapter 6, "STORE and TurboSTORE Command Options," for information on using the `ONLINE` option only, or Chapter 7, "TurboSTORE/iX 7x24 True-Online Backup," for information on using the new `ONLINE=START` and `ONLINE=END` options.

Using Unlabeled vs Labeled Media

`STORE` and `RESTORE` let you store and restore files using either labeled or unlabeled tapes. All TurboSTORE options, except sequential, are supported for both formats and there is no performance penalty for using either format.

The use of labeled tape offers the following advantages:

1. Tapes are self identifying.
2. Automatic Volume Recognition (AVR) will eliminate need for operator replies to media requests.
3. TurboSTORE provides hooks for third party tape management systems.

Maintaining a Backup Library

Establish a library for storing backup media and backup records. Include the following in your library:

- A large stock of new or scratch media to use for backups. (Scratch media are used media containing information that is no longer needed.)
- A record of each tape's (or other media) contents. It should be easy to find the backup media needed to restore the system in the event of a failure.
- Your actual backup media.

Maintaining a Large Stock Of Backup Media

You need to maintain a large and easily accessible stock of available backup media, because you can never be sure ahead of time exactly how many you will need for each backup. The amount of media you will need depends on how much data is being stored, and what STORE and TurboSTORE options, such as data compression, are being used. The amount of data on your system, and whether you are doing a full or partial backup also greatly influence the amount of backup media required.

Using new units of backup media ensures the best possible data integrity. You can, of course, safely reuse backup media if they are in good condition and if they contain outdated information. The information on used backup media is outdated when the purge date on the label has passed.

To help judge whether a used tape is in good condition, you can use the cycle field of the Hewlett-Packard standard sticky tape label to track the number of times a tape has been used. Tapes wear out and become unreliable after being used many times. To determine whether a tape is still usable, consult the manufacturer's recommendations.

When using labeled media, the expiration date in the label can be used to recycle media.

Keeping Backup Records

Each time you create a backup, have STORE display a listing describing what files were stored on each piece of media. STORE's SHOW=OFFLINE option sends this listing, by default, to the line printer. For example:

```
: FILE T;DEV=TAPE
: STORE @.OPERATOR.SYS;*T;SHOW=OFFLINE
```

File these reports in a special filing cabinet or attach them directly to the appropriate tape reel (or other media). Label each item according to the instructions provided in Chapter 3, "Preparing Storage Devices."

Storage of Backup Media

Backup media may be the only copy of current or previous information stored from your system. When you need to restore information from backup media, you should be able to easily find the item you need. Your library enables you to maintain complete copies of your

system's data, and an easy method for locating the data you want to recover.

There are two reasons you should always keep older backup data in addition to your current backup data: in case the most recent copies are bad, and for restoring older files that may have been removed from the system some time ago. In such situations, you can restore files from the most recent intact copy.

To set up a library, perform the following steps:

- Obtain a substantial number of backup media. The exact number you need depends on how much data you will be storing and how long the backup data will be retained.
- Obtain a container such as a vault, a metal cabinet, or rack in which to keep the backup data.
- Establish a control program for maintaining backup data for a specified length of time.
- Establish a cleaning and verifying schedule for the media devices.
- As part of your disaster recovery plan, create a duplicate library in a safe, offsite location. Refer to Chapter 13, "Disaster Recovery," for more information.

3 Preparing Storage Devices

This chapter describes how to prepare your media and devices for backup. The sections in this chapter are:

- Preparation Process
- Preparing Tape and DDS Media
- Preparing Tape and DDS Devices

Preparation Process

The following is a brief list of the preparation steps.

1. You must first determine what type of device(s) to use for your backup. Most of this section deals with storing to magnetic tape or DDS devices. These procedures can also be followed if you are using a StorageTek device. If you want to use disk files, refer to the sections "Storing to Disk Files" in this chapter.
2. If you are using tape or DDS devices, you must next decide if you will use a single device or multiple devices. All versions of STORE and TurboSTORE can store to multiple devices. The possible configurations available are:

Single	Requires only one backup device. When the media (tape or DDS) fills, the device rewinds the media, you remove the media and replace it with new media, and additional files are stored to the new media.
Sequential	Requires more than one backup device. One device is used at a time. When the media (tape or DDS) fills, additional files are stored on the next device in line. You can set this up to toggle between two devices, giving you time to change media.
Parallel	Requires more than one backup device. Files are being stored to more than one device at the same time. When the media is full on one device, you remove the media and replace it with new media, and additional files are stored to the new media.

Parallel Pools Requires more than two backup devices. Files are being stored to more than one device at the same time. This is a combination of Sequential and Parallel. When the media (tape or DDS) fills, additional files are stored on the next device in line. You can set this up to have multiple pairs of toggling devices.

- Decide on the order that each backup device in the pool is to be used.
- Assign each device a position in the backup, identifying each device by its logical device (LDEV) number. Use the chart in Figure 3-1. as a tool to help you assign your device. Sequential devices are numbered 1 through 8. Parallel devices are identified by the letters A through P.

Write the LDEV numbers in the order and configuration that you plan to do the backups. Note that the chart for parallel device pools shows only two sequential devices in each parallel set. If you are using your devices in a different parallel pool configuration, you should create your own chart and identify the position of each LDEV.

Figure 3-1. Multiple-Device Planning Chart

Planning for Multiple Devices Select either Sequential, Parallel, or Parallel Device Pools. Fill in LDEV numbers for each device you plan to use.			
<input type="checkbox"/> Sequential			
Device 1 _____	Device 2 _____	Device 3 _____	Device 4 _____
Device 5 _____	Device 6 _____	Device 7 _____	Device 8 _____
<input type="checkbox"/> Parallel			
Device A _____	Device E _____	Device I _____	Device M _____
Device B _____	Device F _____	Device J _____	Device N _____
Device C _____	Device G _____	Device K _____	Device O _____
Device D _____	Device H _____	Device L _____	Device P _____
<input type="checkbox"/> Parallel Device Pools			
Device A1 _____	Device A2 _____	Device I1 _____	Device I2 _____
Device B1 _____	Device B2 _____	Device J1 _____	Device J2 _____
Device C1 _____	Device C2 _____	Device K1 _____	Device K2 _____
Device D1 _____	Device D2 _____	Device L1 _____	Device L2 _____
Device E1 _____	Device E2 _____	Device M1 _____	Device M2 _____
Device F1 _____	Device F2 _____	Device N1 _____	Device N2 _____
Device G1 _____	Device G2 _____	Device O1 _____	Device O2 _____
Device H1 _____	Device H2 _____	Device P1 _____	Device P2 _____

LG200021_040

Selecting the Backup Method

The best way to use your backup devices depends upon the characteristics of the files that you are storing, your needs for compatibility with other versions of MPE, and your system configuration.

The characteristics of the set of files help determine choices (or options) for storing them. Generally speaking, the more files being stored, the more it is to your advantage to use multiple devices. If you are storing only a few files of moderate size, one backup device may be sufficient. If you are storing many files (for example, all the files in an account, or a smaller number of very large files), using a multiple-device backup is probably to your advantage.

Your needs for compatibility with other versions of MPE limit your choices. If you need to produce `STORE` backups that are compatible with MPE V/E using the `TRANSPORT` parameter, you can only store files to a single device.

If you are not creating tapes for use by MPE/V, then your system configuration is the most important determinant of which multiple-device backup method to use. Use the following guidelines to select a method:

Which and How Many Backup Devices Do You Have?

Which devices you may use for a multiple-device backup are described in this chapter. Obviously, to use sequential or parallel devices, you must have at least two backup devices. If you have two or three devices, you may use either parallel or sequential devices. If you have three or more devices, you may use two or more parallel device pools.

If you want to use magneto-optical devices or disk files, refer to the sections "Preparing Magneto-Optical (MO) Disks" and "Storing to Disk Files" later in this chapter. If you are using StorageTek devices, most of the guidelines for tape and DDS devices will apply.

If you are using StorageTek devices, you should follow the directions for tape and DDS devices. However, note that due to the higher throughput of these devices, they should be configured on your system for maximum performance. Refer to the Configuration Guide (put real name there) for information on the proper configuration for 3480 devices.

Selecting Backup Device Order

Once you have determined how many devices you are going to use for your backup, you should determine the order of each device in the backup. Assign each device a position based on its logical device (LDEV) number.

Use the chart in Figure 3-1. to help you assign your devices. Select one of the methods in the chart, and then write in the LDEV numbers of each device you plan to use in the order that you plan to use it. For example, if you intend to use two parallel devices, write in the LDEV number of the first device next to Device A and the LDEV number of the second device next to Device B in the Parallel section of the chart. It may be helpful for you to photocopy the chart and use it as a guide whenever you use a multiple-device method to store files.

Assigning Device File Names

When using multiple backup devices, you have the choice of entering multiple `FILE` commands to assign a different filename to each device or a single `FILE` command.

Assigning each device a different filename makes it easier to keep track of each device's position in the backup media. When you receive media requests, you know exactly to which device the request applies. You may use LDEV numbers or a common device class in the `FILE` command. When you use a device class to name backup devices, `STORE` asks you to assign a logical device to the filename with a media request. Refer to Chapter 8, "Handling Media Requests" for more information. You must use LDEV numbers to name devices if the devices are to be automatically allocated.

Assign unique filenames based on each device's position in the backup. For example, if you plan to use four parallel devices, you might issue the following set of `FILE` commands:

```
: FILE BACKA ;DEV=7
: FILE BACKB ;DEV=8
: FILE BACKC ;DEV=9
: FILE BACKD ;DEV=11
```

Label sequential devices in order. For example, if you are using three sequential devices, you might issue the following `FILE` commands:

```
: FILE BACK1 ;DEV=TAPE
: FILE BACK2 ;DEV=TAPE
: FILE BACK3 ;DEV=TAPE
```

Label parallel device pools both in parallel and sequential order. For example:

```
: FILE BACKA1 ;DEV=7
: FILE BACKA2 ;DEV=8
: FILE BACKB1 ;DEV=9
: FILE BACKB2 ;DEV=11
```

If you issue a single `FILE` command when using multiple backup devices, you assign a common filename and common device class to all of your backup devices. For example:

```
: FILE T ;DEV=TAPE
```

You use the common filename, `T` in the example above, to refer to each device in your `STORE` command. In a media request, the system asks you to assign an LDEV number to `T` for each device. If you have four backup devices, you see four identical media requests. You assign a different LDEV number to each media request based upon the order in which you receive the requests. Refer to Chapter 8, "Handling Media Requests," for more information.

Your `FILE` command gives the backup device a filename and identifies the device. To store files to only one device, you backreference that device in your `STORE` command, as described previously. For example:

```
: FILE T ;DEV=TAPE
: STORE @.OPERATOR.SYS ;*T
```

However, you may choose to employ multiple device configurations at times when you have a large set of files to store and you want to speed up the process. If so, you probably want to use the `FILE` command to give each device a filename that reminds you of its position. Then use the `STORESET` parameter in a `STORE` command to reference the multiple devices. For example, to store files to three sequential devices, use the following set of commands:

```
:FILE SEQ1;DEV=7
:FILE SEQ2;DEV=8
:FILE SEQ3;DEV=9
:STORE @.@@;STORESET=( *SEQ1 , *SEQ2 , *SEQ3 )
```

The above `STORE` command copies all system and user files to three sequential devices. The parentheses surrounding backreferenced file names group the devices into a device pool.

NOTE You do not backreference the filename of a backup device immediately after the name(s) of the file(s) you are storing when you use the `STORESET` parameter. Rather, you must use a placeholder for the missing parameter normally supplied when using a single backup device. The placeholder is provided by the use of `;;` in the `STORE` command.

If you gave the backup devices a single filename such as:

```
:FILE T;DEV=TAPE
```

enter the following `STORE` command to copy files to a set of four sequential backup devices:

```
:STORE @.@@;STORESET=( *T , *T , *T , *T )
```

To store files to a set of four parallel devices, you might use the following set of commands:

```
:FILE PARA;DEV=7
:FILE PARB;DEV=8
:FILE PARC;DEV=9
:FILE PARD;DEV=11
:STORE @.OPERATOR.SYS; ;&
:STORESET=( *PARA ) , ( *PARB ) , ( *PARC ) , ( *PARD )
```

If you gave the backup devices a single filename such as:

```
:FILE T;DEV=TAPE
```

enter the following `STORE` command to copy files to a set of four parallel backup devices:

```
:STORE @.@@;STORESET=( *T ) , ( *T ) , ( *T ) , ( *T )
```

You describe parallel devices in the `STORESET` parameter by enclosing each backreferenced device filename within parentheses. Use commas to separate devices.

To store files to parallel device pools, use parentheses to group the devices into pools. Separate the devices in the pool from one another with commas. Commas also separate each device pool. For example, the following commands copy files to two parallel device pools. Each device pool contains two sequential devices.

In the following example, STORE uses a total of four devices divided into two device pools.

```
:FILE PPA1;DEV=7
:FILE PPA2;DEV=8
:FILE PPB1;DEV=9
:FILE PPB2;DEV=11
:STORE @.OPERATOR.SYS;;STORESET=( *PPA1 , *PPA2 ) , ( *PPB1 , *PPB2 )
```

If you gave the backup devices a single filename such as:

```
:FILE T;DEV=TAPE
```

enter the following STORE command to copy files to a set of parallel device pools:

```
:STORE @.@.@@; ;STORESET=( *T , *T ) , ( *T , *T )
```

NOTE Up to 16 parallel devices may be specified for one backup. With a serial device pool, up to 8 devices may be specified.

When using a multiple-processor machine, you can see increased backup performance when doing parallel backups, since each parallel device gets its own process. A general rule for doing multi-processor parallel backups is to use one less parallel set than you have processors. For example, on a 4-way system, perform a 3-device parallel backup.

Preparing Tapes and DDSs

This section discusses preparing tape and DDS devices. Much of the information can also be applied to storing to StorageTek devices, although some implementation details may differ. If you are using disk devices for your backup, refer to the sections "Storing to Disk Files" for more information on those devices.

You must select the backup media onto which you are going to store your files. If you have followed the instructions for setting up a backup library in Chapter 2, "Managing Backups and Backup Media," you should have a set of new or scratch backup media or both available for storing files.

To prepare your backup media, perform the following steps:

1. Estimate the number of blank media you need, and take them from your library.
2. Attach a label to each unit of backup media. (A standard Hewlett-Packard tape label has several headings for different types of important information. If your label does not have such headings, you may want to create your own headings including the information which is described next.)

3. Label the first unit of backup media.

Single or Sequential

Label the first unit Volume 1 of n . Replace n with the total number of units of backup media--if you use more or fewer tapes than you planned, you may have to change this number when you finish storing the files.

Label the second unit of backup media Volume 2 of n . Label the third unit of backup media Volume 3 of n , and so on.

Parallel or Parallel Pools

Label each piece of media with both the reel number **and** the set number. Label the first media of each parallel device or device pool as

Volume 1 of "x", Set 1 of "y"

where "x" is the total number of media allocated to the particular parallel set, and "y" is the total number of parallel sets. If you use more or fewer units of backup media than you planned, you may have to change the final number of volumes when you finish.

For example, if you plan to use four units of backup media with two parallel devices, label:

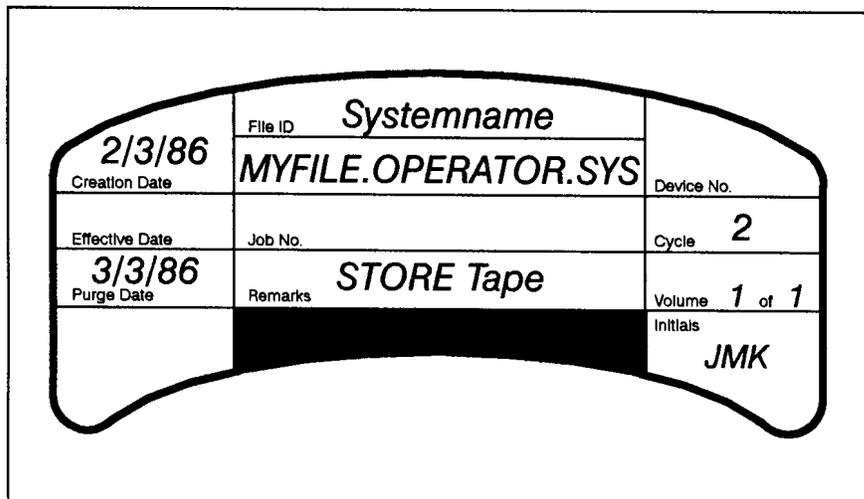
- a. First media as Volume 1 of 2, Set 1 of 2.
- b. Second media from the first parallel device as Volume 2 of 2, Set 1 of 2.
- c. First media from the second parallel device as Volume 1 of 2, Set 2 of 2.
- d. Second media from the second parallel set as Volume 2 of 2, Set 2 of 2.

4. In addition to the media number, write the following information on each media's label:

- The current date which serves as a reminder of when the backup was created.
- The name of your computer system.
- A list of the files you intend to store, or if the tape is from a system backup, write either "FULL BACKUP" or "PARTIAL BACKUP".
- If you plan to use interleave format, write "INTERLEAVE". (Interleave format is described in Chapter 5, "Storing Files.")
- Any other important information about the backup. Since it is a STORE backup, write "STORE".
- Your name or initials.
- The date after which the information on the backup is obsolete, called the Purge Date. If you keep STORE backups indefinitely, make a note to that effect.
- The number of times the backup media has been used, under the heading "Cycle".

Figure 3-2. shows an example of a correctly labeled backup:

Figure 3-2. A Correctly Labeled Tape

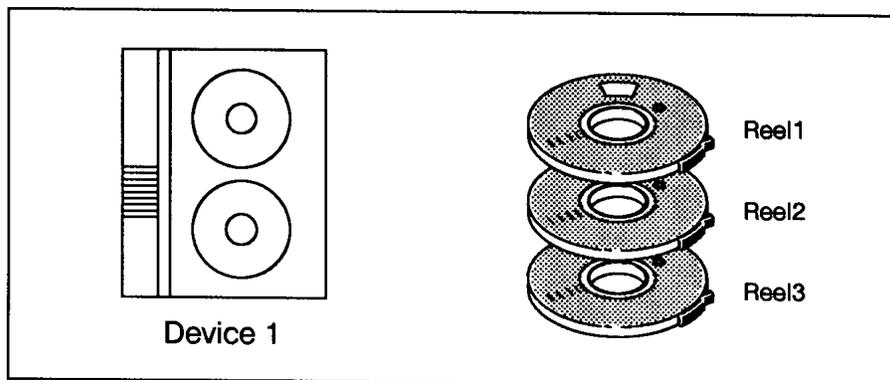


LG200021_022

5. Stack all your media (tapes or DDSs).

Single If you plan to use a single device, stack all of your backups, label side up, in order, with the lowest number on top, near the device, as shown in Figure 3-3..

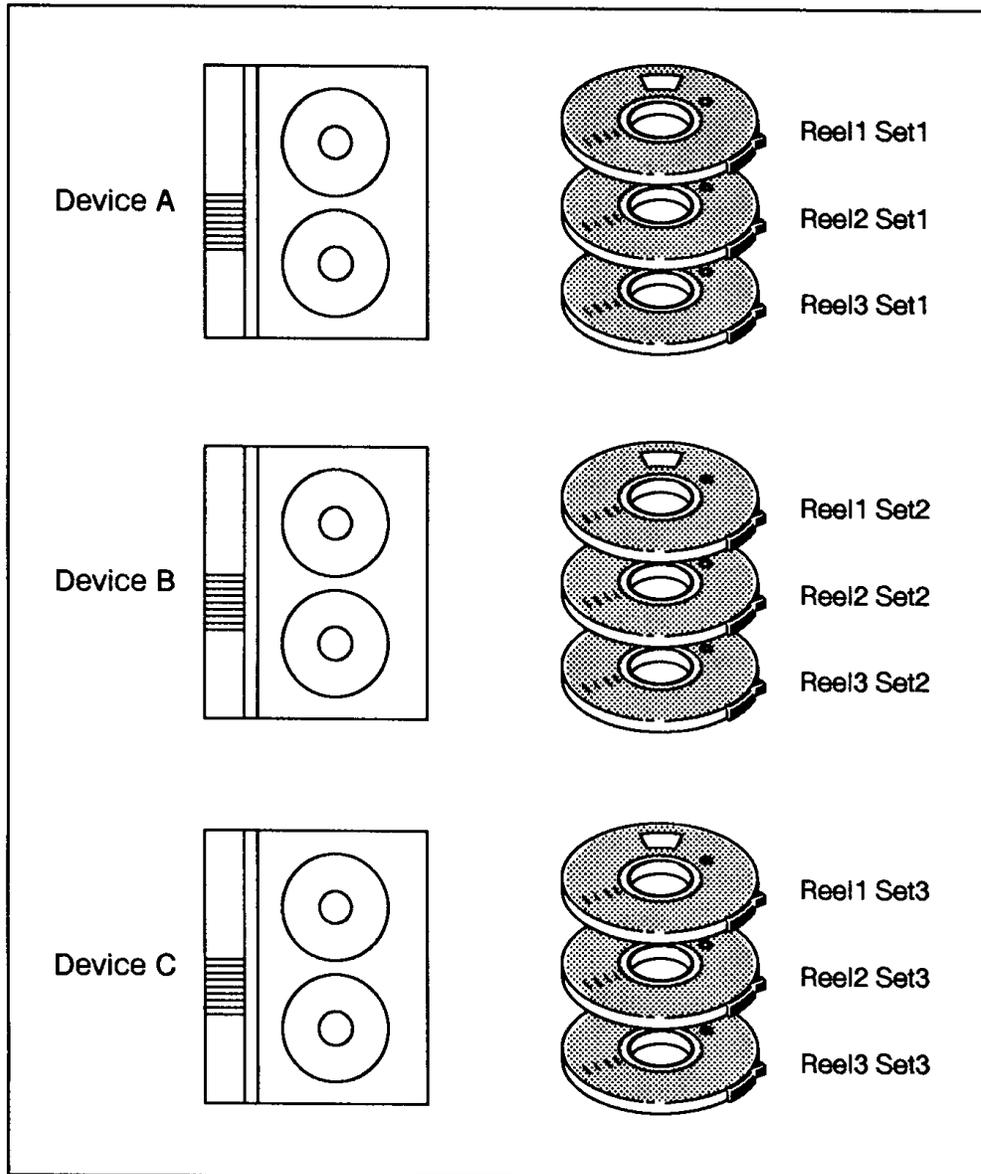
Figure 3-3. Preparing Tapes for a Single Device



LG200021_007a

Parallel If you are using parallel devices, stack all of your backups, label side up, in order with the lowest number on top, near the corresponding device. For example, place all the backups with numbers ending with Set1 near your first parallel device. Put media Reel1 Set1 on top of the stack. Figure 3-4. shows how to prepare backup medias for parallel devices.

Figure 3-4. Preparing Tapes for Parallel Devices



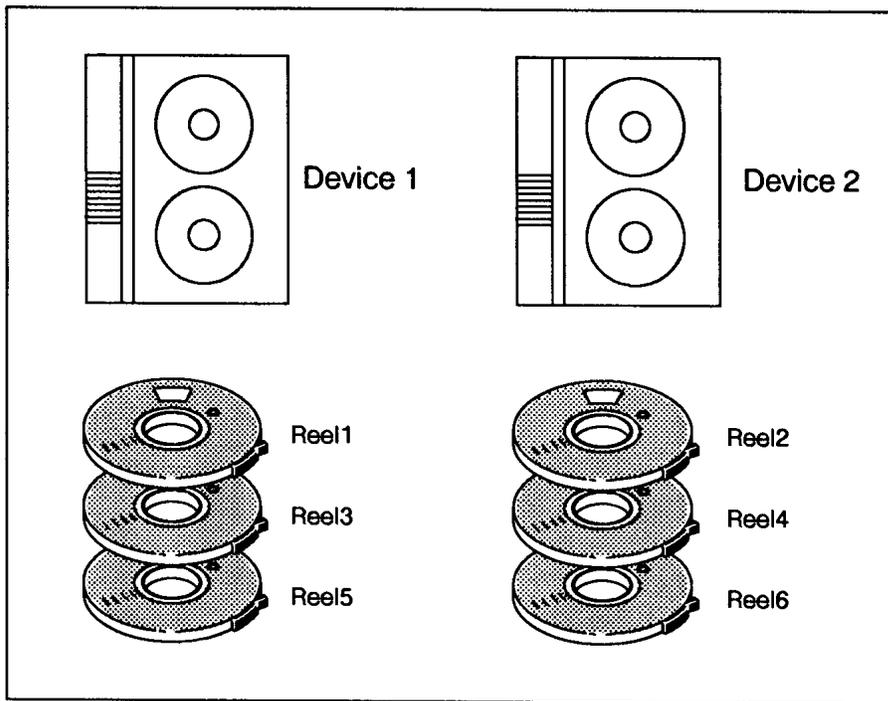
LG200021_009a

6. Divide your media among the devices.

Sequential

If you plan to use sequential devices, divide your media among the devices in order. For example, if you have two sequential devices and six media, stack media 1, 3 and 5 near device 1, and stack media 2, 4 and 6 near device 2. Stack all media label side up and make sure that the lowest-numbered backup is on the top of each stack. Figure 3-5. shows an example.

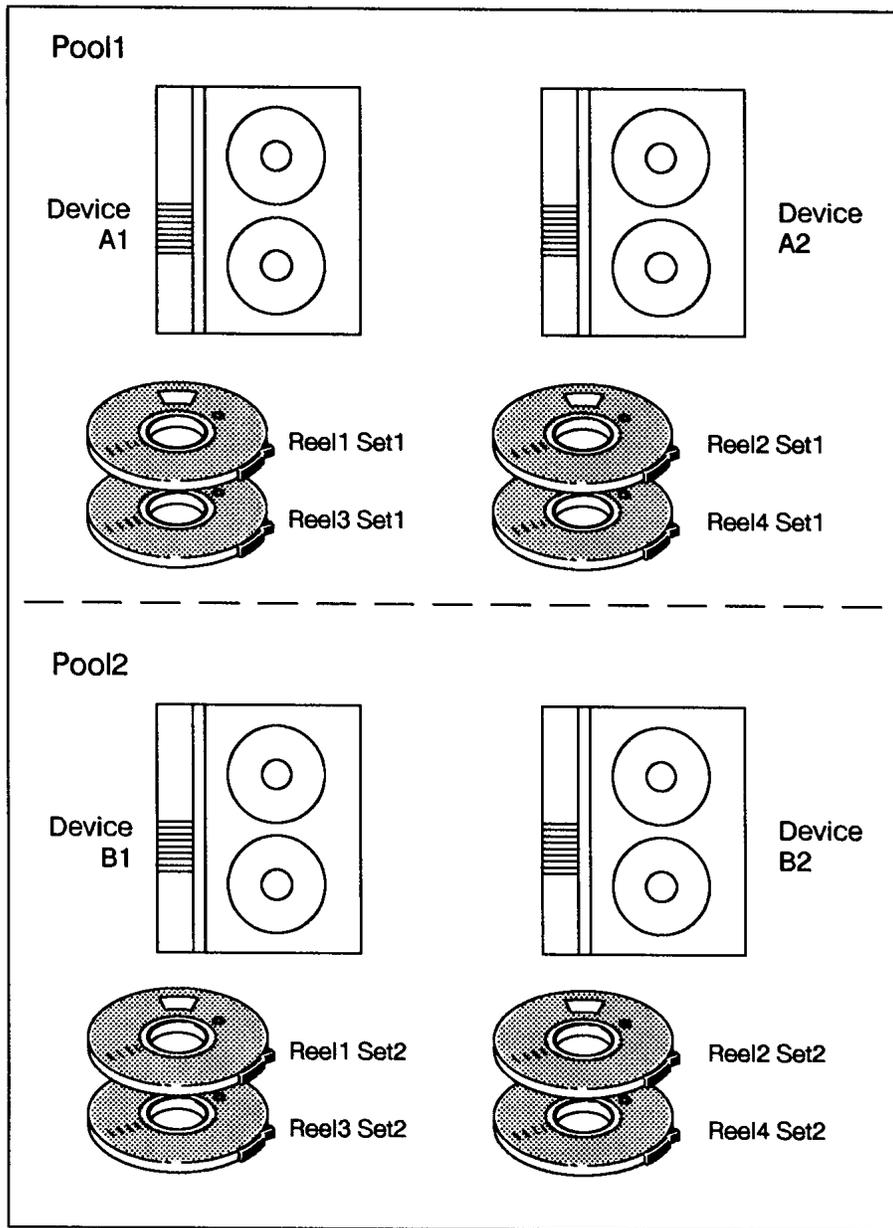
Figure 3-5. Preparing Tapes for Sequential Devices



LG200021_008a

Parallel Pool If you are using parallel device pools, divide your media among the devices in each device pool, placing them in order with the label side up. For example, if the first device pool has two sequential devices and four media, stack media Reel1 Set1 and Reel3 Set1 near device A1, and place media Reel2 Set1 and Reel4 Set1 near device A2. Figure 3-6. shows media prepared for storing files to parallel device pools.

Figure 3-6. Preparing Tapes for Parallel Device Pools



LG200021_010b

Preparing Tape and DDS Devices

Prepare your backup devices by mounting the first media in the stack near each device onto that device. Follow the instructions in Chapter 3, "Preparing Storage Devices," for mounting backup media.

After mounting the media, check the console for a message telling you that the volume has been mounted. For example:

```
14:47/31/VOL (Unlabeled) mounted on LDEV#8
```

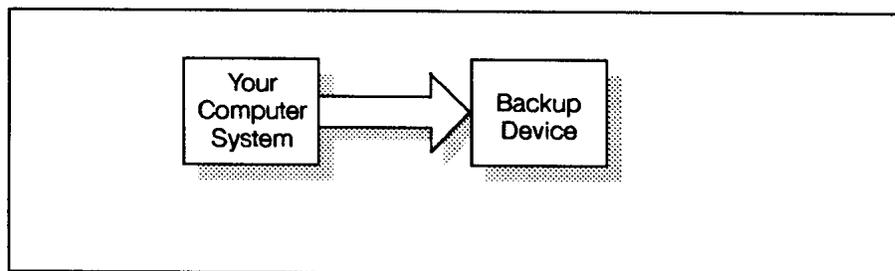
Note the backup drive's logical device (LDEV) number. (The LDEV number is 8 in the example above.) You may need it to answer a backup request later.

Single Tape Devices

When you store to a single backup device, `STORE` only writes files to one device at a time. When the backup media is full, it automatically rewinds then prompts you to load new backup media. When the second unit of backup media is loaded, `STORE` continues writing files until all of the qualified files are stored or another unit of backup media is needed. If you are storing many files, or files that are particularly large, you may need individually load several units of backup media onto the single device.

Figure 3-7. illustrates a single-device backup.

Figure 3-7. Single-Device Backup



LG200021_001

The `FILE` command assigns a filename to a backup device according to its logical device number or its device class. For example, the `FILE` command below assigns the filename `T` to the device with LDEV 8:

```
: FILE T;DEV=8
```

To store files to only one device, backreference that device in your `STORE` command. For example:

```
: FILE T;DEV=TAPE
```

```
: STORE @.OPERATOR.SYS;*T
```

Using one or more `FILE` commands is recommended, especially when your system is

configured to automatically allocate devices (autoreply) or when you are using multiple backup devices.

NOTE If auto reply is not enabled on the backup device designated, a reply from the console is necessary to allocate that device for a `STORE` process. See Chapter 8, "Handling Media Requests," for more information.

If you are using labeled tapes for your backup, you may or may not have to respond to console requests, depending on how your system is configured. Refer to the "Storing to Labeled Tapes" section in Chapter 5 for more information on using labeled tapes.

The `FILE` command below assigns the filename `BACKUP` to the device with the device class name `TAPE`.

```
:FILE BACKUP;DEV=TAPE
```

Using the backup device's logical device number is a more exact way to assign a filename to a particular device. There is only one device defined as `LDEV 8` on your system. However, several devices may be assigned the device class `TAPE`. When you use device classes to name backup devices, `STORE` asks you to assign a logical device number to each filename with a backup request.

FILE Commands for Automatically Allocated Devices

In order to have devices automatically allocated to the backup, you must assign a different filename for each backup device allocated; use an `LDEV` number or a unique device class for each `FILE` command you use, one for each device.

Refer to Chapter 8, "Handling Media Requests," for a description of automatically allocated devices.

Multiple Tape Devices

All `STORE` and `TurboSTORE` products allow the use of multiple devices for backup. Multiple devices are identified as a pool of available devices. When you store to multiple devices, the devices must have similar characteristics; different device types are not allowed. The device characteristics that must be identical are:

- Devices must be all local or all remote.
- All devices must have the same density.
- All devices must have the same record size.
- All backups must be either labeled or unlabeled; there cannot be a mix.
- Different types of backup devices cannot be used together in one backup. A backup must be performed with only one type of backup device at a time. Note that different models of DDS drives (i.e., SCSI or HPIB) can be mixed. Different models of magnetic tape drives (7980, 7980XC, etc.) can also be used together.

Multiple devices can be used as:

- Sequential devices.
- Parallel devices.
- Parallel device pools.

Using Tape Devices Sequentially

With sequential backup devices, when the first backup on the first backup device is filled, it immediately begins writing files to a second backup on the second backup device and the first backup rewinds. When the first backup is rewound, you can unload it and prepare the drive with a new backup. You do not lose time while the first backup is rewinding nor while you are loading a new backup.

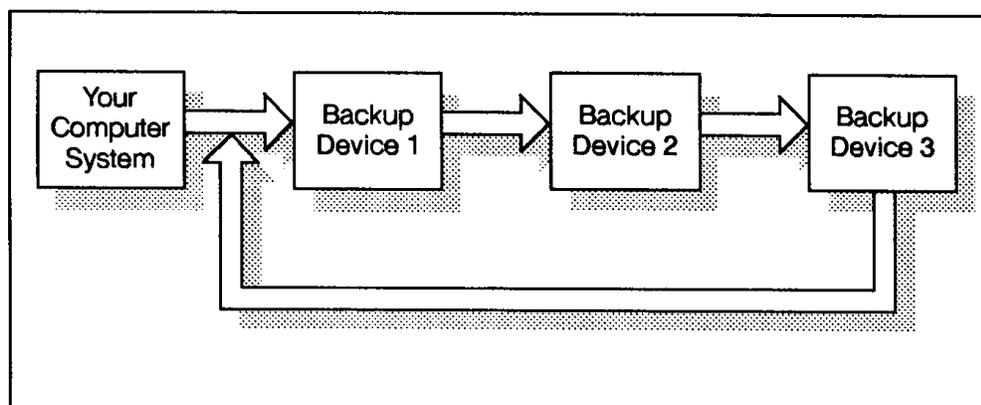
When the backup on the last sequential device is filled, the process returns to the first device. It continues cycling through all the devices in the device pool until all of the specified files are stored.

The maximum number of serial magnetic tape or DDS devices allowed is 8. However, you generally do not need to use more than two devices in serial. You get better backup performance if you use multiple devices in parallel. However, using two devices in serial can be useful for some unattended backup functions. For example, if your backup usually fits on one media, then using two devices in serial allows your backup to finish unattended, even if your data grows and one night overflows one piece of media.

Serial devices can also help a backup to continue unattended in the event of a media write error. If STORE encounters an error writing to a backup media, it marks the current media bad and unloads it. It then rewrites the backup data to the next mounted serial device, and the backup continues uninterrupted. This gives you enough time to mount another media on the other drive, if needed.

Figure 3-8. illustrates a sequential device backup.

Figure 3-8. Sequential Backup



LG200021_002

Use the `FILE` command to give each device a filename that reminds you of its position. Then use the `STORESET` parameter in a `STORE` command to reference the multiple devices. For example, to store files to three sequential devices, use the following set of commands:

```
:FILE SEQ1;DEV=7
:FILE SEQ2;DEV=8
:FILE SEQ3;DEV=9
:STORE @. @. @.;;STORESET=( *SEQ1 , *SEQ2 , *SEQ3 )
```

The above `STORE` command copies all system and user files to three sequential devices. The parentheses surrounding backreferenced filenames group the devices into a device pool.

NOTE You do not backreference the filename of a backup device immediately after the name(s) of the file(s) you are storing when you use the `STORESET` parameter. Instead, you must use a placeholder for the missing parameter normally supplied when using a single backup device. The placeholder is provided by the use of ";" in the `STORE` command.

If you gave the backup devices a single filename such as:

```
:FILE T;DEV=TAPE
```

enter the following `STORE` command to copy files to a set of three sequential backup devices:

```
:STORE @. @. @.;;STORESET=( *T , *T , *T )
```

NOTE Only unlabeled media can be used for a sequential device backup. A maximum of 8 devices can be specified in each serial device pool.

Using Tape Devices in Parallel

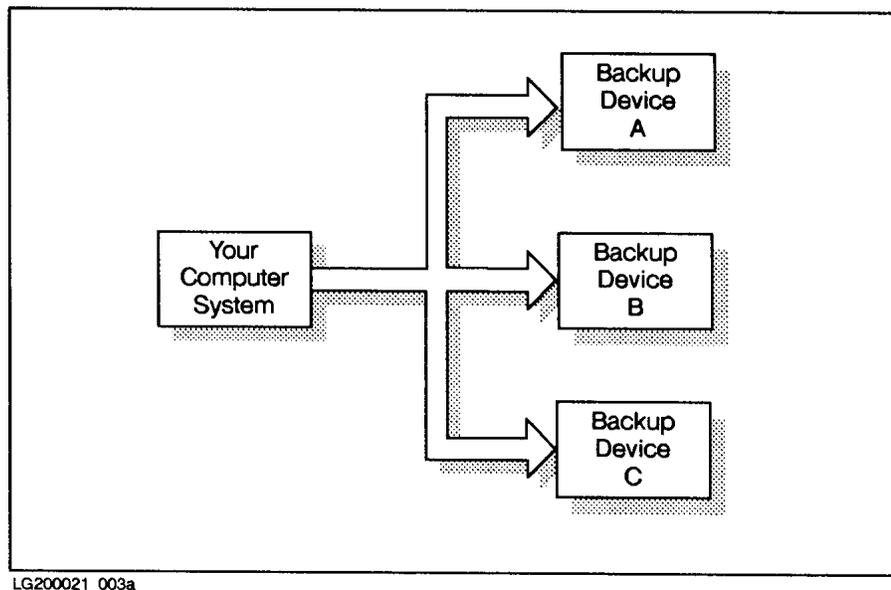
With parallel devices, you have the capability of reducing backup time by storing different files on different devices at the same time. When you store files to a set of parallel devices, the files are grouped before the store begins and then copied to multiple devices simultaneously. For example, when you have three parallel devices, your files are partitioned into three subsets. Files in the first subset are copied to the first backup drive, the files in the second subset are copied to the second backup drive, and so on.

When a piece of media is filled, you are prompted to mount the next piece of media on that device.

Depending upon the characteristics of your files and your backup devices, storing files to parallel devices usually saves more time than using sequential devices.

Figure 3-9. illustrates a parallel device backup.

Figure 3-9. Parallel Backup



Currently, all STORE and TurboSTORE products support a maximum of 16 parallel backup devices. The actual speed of performance depends on the system, disk configuration and the backup devices used.

When using a multiple-processor machine, you can see increased backup performance when doing parallel backups. Each parallel device or device pool is controlled by a separate process. Since each process can be run simultaneously on different processors, this will improve your backup performance.

A general rule for doing multi-processor parallel backups is to use one less parallel set than you have processors. For example, on a 4-way system, perform a 3-device parallel backup. This is recommended because STORE uses two additional processes to control the overall backup.

NOTE Depending on your system configuration and bandwidth, you should still see increased backup performance using parallel devices, even if you use more parallel devices than you have processors.

To store files to a set of four parallel devices, the following set of commands can be used:

```
: FILE PARA ; DEV=7  
: FILE PARB ; DEV=8  
: FILE PARC ; DEV=9  
: FILE PARD ; DEV=11  
: STORE @.OPERATOR.SYS ; ; &  
: STORESET= ( *PARA ) , ( *PARB ) , ( *PARC ) , ( *PARD )
```

If you gave the backup devices a single filename such as:

```
:FILE T;DEV=TAPE
```

enter the following STORE command to copy files to a set of four parallel backup devices:

```
:STORE @.@@;STORESET=(*T),(*T),(*T),(*T)
```

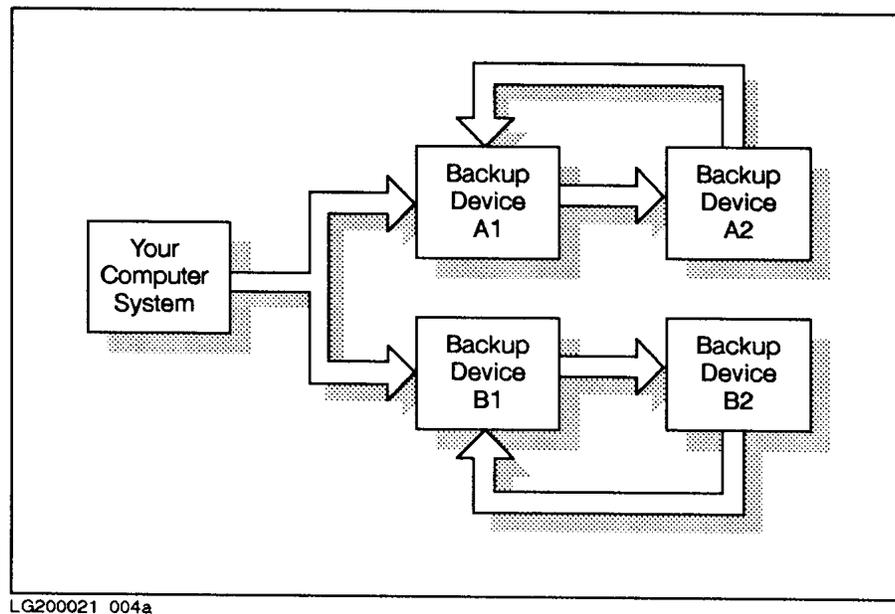
Describe parallel devices in the STORESET parameter by enclosing each backreferenced device filename within parentheses. Use commas to separate devices.

Using Parallel Device Pools

Parallel device pools are multiple sets of sequential devices used in parallel. Using parallel device pools is often the most efficient way to store a large set of files when you have several backup devices. Parallel device pools give you the advantages of both parallel and sequential devices. You copy file subsets to different devices simultaneously, and can have another sequential device ready when the media is filled in any device pool.

Figure 3-10. illustrates a parallel device pool backup.

Figure 3-10. Parallel Device Pools



Up to eight serial backup devices are currently supported within a parallel device pool. Maximum performance is obtained by using three or four parallel sets, each consisting of two sequential devices.

To store files to parallel device pools, use parentheses to group the devices into pools. Separate the devices in the pool from one another with commas. Commas also separate each device pool. For example, the following commands copy files to two parallel device pools. Each device pool contains two sequential devices. In the following example, STORE

uses a total of four devices divided into two device pools.

```
:FILE PPA1;DEV=7
:FILE PPA2;DEV=8
:FILE PPB1;DEV=9
:FILE PPB2;DEV=11
:STORE @.OPERATOR.SYS;;STORESET=( *PPA1,*PPA2 ), ( *PPB1,*PPB2 )
```

If you gave the backup devices a single filename such as:

```
:FILE T;DEV=TAPE
```

enter the following `STORE` command to copy files to a set of parallel device pools:

```
:STORE @. @. @.;;STORESET=( *T,*T ), ( *T,*T )
```

Using \$NULL as a STORE Device

The devices specified for a backup can also refer to `$NULL`. This essentially does the backup to the "bit bucket" -- no data is stored to any backup media. This type of backup can be done by specifying a file equation prior to performing the backup. Instead of specifying a tape device, specify `$NULL`. For example:

```
FILE N=$NULL
```

Doing so causes `STORE` to use `$NULL` as its output device. The command line is parsed, the input fileset scanned, all files are read from disk, and all output is displayed as normal. The only real difference is that the data is not actually being written to an output device.

This feature can be useful when trying to determine which files will be stored by `STORE`, as a trial run. It can also be used, with **EXTREME** caution, with the `;PURGE` option, to remove all specified files.

When using this feature, keep in mind that your files are **NOT** being stored to any device, so the data has not been saved anywhere.

Using Remote Devices for Backups

`STORE` and `TurboSTORE` now support remote backup devices for storing and restoring files. This feature, which was previously only available using `CMSTORE`, allows you to specify a backup device, or devices, on a remote system for your backups. A connection to the remote system must exist before the `STORE` is invoked. Once the remote connection is established, a file equation should be created to point to the remote device. For example, to store to a tape drive on the remote system `BROWNY`, issue the following `FILE` command:

```
FILE REMOTE;DEV=BROWNY#TAPE
```

This file equation can then be used as the device file argument to `STORE`, `RESTORE`, or `VSTORE`, or can be used with the `STORESET` or `RESTORESET` options. For example, using the file equation defined above, the following `STORE` command will back up files to a remote tape drive on `BROWNY`:

```
STORE @. @. @.; *REMOTE;SHOW
```

Although the initial console device request will appear on the console of the remote system, all further `STORE` console messages will be sent to the *local* system's console. These

messages will also be sent to the backup's \$STDLIST.

NOTE Remote backup is currently only supported for *unlabeled* backups. Labeled backups cannot be accessed remotely by either STORE or CMSTORE.

Storing to Disk Files

Using TurboSTORE/iX 7x24 True-Online Backup (product number HPB5152AA), you have the ability to store files to disk files. This feature allows you to quickly back up filesets. It also allows you to create archive files that maintain all MPE file attributes. Since the disk backup file created is a flat, binary disk file, it can be moved to different systems, or even different operating systems, and then RESTORE can be used to extract the archived files. Disk files created by TurboSTORE can only be created and restored on systems running MPE/iX 5.5 or later. They cannot be restored on a pre-5.5 system.

File equations are used to tell TurboSTORE to store to a disk file. Instead of setting up the store device file equation to reference a tape device, set up the file equation to point to a disk file. For example:

```
: FILE MYBACKUP ;DEV=DISC
OR
: FILE MYBACKUP=FULLBACK.BACKUPS.SYS ;DEV=DISC
```

When specifying the file equation, ;DEV=DISC MUST be specified in order to Store to a disk file. The target of the file equation can point to any MPE or HFS filename. Any other information specified in the file equation will be ignored by TurboSTORE, so only the filename and DEV=DISC should be used.

It is strongly recommended that you do NOT build the target disk file before invoking STORE. TurboSTORE will create the target disk file as a binary file, with a default record size of 256 bytes. The file will also be created with the file code 2501. The file system mnemonic for this is STORE. For example, here is a listing of a TurboSTORE disk backup file created by STORE:

```
:LISTFILE MYBACKUP,2

FILENAME  CODE  -----LOGICAL RECORD-----  ----SPACE----
          SIZE  TYP      EOF      LIMIT R/B  SECTORS #X MX
MYBACKUP  STORE  128W  FB      8736   16776959  1    10272  8  *
```

Although it is not recommended, you can create the target disk files before starting the backup. Only normal, fixed, or undefined length record files can be specified. STORE opens the file for user-mapped write access, so any restrictions that apply for user-mapped access apply for store disk files. Store will NOT change the file code of an existing disk file, so if you want the file to have the proper file code, specify it when you build the file.

If STORE fills up the first disk file specified for the backup, it creates as many additional disk files as needed, or uses existing disk files. They will be built with the same default file characteristics as the first disk file. The naming convention used for additional files is to append the reel number to the end of the first disk filename. The resulting name will be an HFS-syntax name.

For example, if STORE needed three disk files to store all files, they would be named:

/SYS/MYBACKUP/STORDISC

/SYS/MYBACKUP/STORDISC.2

/SYS/MYBACKUP/STORDISC.3

When STORE moves on to another file, be sure enough disk space exists before doing the backup.

NOTE The disk files can be restored by the MPE/iX 5.5 Limited Release version of RESTORE or later. RESTORE cannot restore disk files on pre-MPE/iX 5.5 systems.

4 Preparing the System

This chapter describes the procedures for notifying and preparing all the interactive items of a system for a system level backup. The sections are:

- Notifying Users of Pending Backup
- Preventing New System Activity
- Sending Warning Messages
- Replying to Console Requests
- Temporarily Suspending Jobs
- Sending Final Warning Messages
- Aborting Any Remaining Sessions

For non-online backups, it is recommended that system backups be done when no one else is using the system. This ensures that all files are backed up, since files open for write access will not be stored.

For online backups, TurboSTORE/iX 7x24 True-Online Backup has two methods:

- Performing an *online backup* (using the `ONLINE` option only), where users must be logged off the system when the backup is started. Users can get back on the system after the backup has started.
- Performing a *7x24 true-online backup* (using the `ONLINE` option with `START` or `END`), where you never have to ask users to log off of the system when doing a backup. They can continue to access their files and applications throughout the backup. They are available and can be modified while the backup is occurring. Refer to Chapter 7, "TurboSTORE/iX 7x24 True-Online Backup," for more information.

The procedures in this chapter are recommended primarily for doing non-online backups. Many of these steps do not need to be performed when doing an online or 7x24 true-online backup.

NOTE Whenever the notes refer to an *online backup*, this means a backup being performed with the `ONLINE` option. Where the notes refer to a *7x24 true-online backup*, this means a backup being performed with either the `ONLINE=START` option or the `ONLINE=END` option.

To prevent users from accessing the system during a system backup, you must close all files and log all users off of the system before you begin. It is also a good idea to announce backups in the `WELCOME` message. To prepare for a backup, follow the steps below.

Notifying Users of Pending Backup

About 15 minutes before you plan to begin the backup, use the `TELL` command to send a message asking all users to finish their work and to log off the system. Enter the following command substituting the actual time that you plan to begin the backup for `hh:mm`.

```
:TELL @S; BACKUP BEGINS AT hh:mm. PLEASE LOGOFF; &  
:TELL YOUR QUIET NEIGHBORS
```

(A `QUIET` neighbor is someone who has specified not to be disturbed with terminal messages.)

NOTE For an online backup, send this message to everyone logged on to the system, since users should access files being stored at the beginning of an online backup.

For a 7x24 true-online backup, this message does not need to be sent to users. Users can continue to access files throughout the entire backup. However, you may want to send a message to all users informing them that a 7x24 true-online backup is starting.

Preventing New System Activity

You control the number of executing jobs and sessions with two commands, `LIMIT` and `JOBFENCE`. The `LIMIT` command controls the number of jobs or sessions that execute concurrently. (The limits you set with the `LIMIT` command cannot exceed the maximum job/session limit determined by your system configuration.)

The `JOBFENCE` command lets you choose the priority jobfence--a number between 0 and 14 (inclusive). Each session and job logs on to the system with an input priority between 1 and 13 (inclusive); the default input priority for standard MPE users is 8. In order for a session or job to execute, its input priority must exceed the current jobfence.

When you set the jobfence high, jobs with an input priority less than or equal to the jobfence enter the `WAIT` state. MPE XL refuses attempts to initiate a session with an input priority that does not exceed the current job fence. Users see the following message:

```
CAN'T INITIATE NEW SESSIONS NOW.
```

If you set the jobfence to its highest possible value, 14, only users assigned system manager (SM) or system supervisor (OP) capability are able to log on. (System managers and system supervisors are the only users allowed to log on with the `HIPRI` parameter of the `HELLO` and `JOB` commands.) Setting the jobfence to 14 prevents all other users from gaining access to the system during a system backup. Enter:

```
:LIMIT 0,0  
:JOBFENCE 14
```

To check the new job/sessions limits and the system jobfence, enter:

```
:SHOWJOB STATUS
```

The last line of the message on your console should now exactly match the example shown below:

```
JOBFENCE= 14; JLIMIT= 0; SLIMIT= 0
```

If you see anything else, you made a mistake entering the `LIMIT` and `JOBFENCE` commands. Follow the instructions again.

NOTE For an online backup, make sure that no one logs onto the system while the backup is starting. Once TurboSTORE informs you that files can be modified, via a console message, you can allow users to log on again.

For a 7x24 true-online backup, you do not have to prevent new system activity. All normal system activity can continue during the backup.

Even though jobs and sessions can continue to run during both an online and 7x24 true-online backup, it is not advised that you run processor or disk intensive jobs during a backup. Doing so will degrade both the backup's and the job's performance. If possible, try to schedule large jobs and backups at different times.

Sending Warning Messages

About 10 minutes before you plan to begin the backup, use the `WARN` command to send a message to all users asking them to log off the system.

Unlike the `TELL` command, the `WARN` command interrupts users working in `QUIET` mode so that you can be sure they know about the impending backup. Enter the command below, substituting the time you intend to begin the backup for `hh:mm`.

```
:WARN @S; BACKUP WILL BEGIN AT hh:mm. PLEASE LOG OFF
```

Now that everyone has been warned, find out who is still using the system. Enter:

```
:SHOWJOB JOB=@S
```

The system responds by listing all sessions. Since you have already told users about the backup once, and warned them again, most users should have ended their sessions. Periodically use the `SHOWJOB` command to see who is still logged on to the system.

NOTE For an online backup, send these messages to users. However, you could change the text of the message to let users know that they can log on after a certain amount of time has passed. The amount of time users have to remain logged off the system depends on how much data is being stored. It can range anywhere from 1 to 30 minutes.

For a 7x24 true-online backup, you can send a warning message to users notifying them that the true-online backup is about to begin. However, you do not have to tell them to log off.

Replying to Console Requests

Before continuing with the backup, check whether there are any console requests pending. Enter:

:RECALL

If the system responds with `NO REPLIES PENDING`, skip to the "Temporarily Suspend Jobs" step below. If there are some console requests that you have not yet answered, respond to them now. Refer to Chapter 8, "Handling Media Requests," for information on answering media requests. If you give users access to a device at this point, you may have to wait until they finish their work before you proceed with the backup.

NOTE For an online backup or a 7x24 true-online backup, you may still want to make sure there are no outstanding console requests. You do not want user jobs competing for backup devices when trying to launch your backup. However, if you are performing a 7x24 true-online backup, and know that the devices being requested will **not** be used for the backup, you can reply to the console requests, and let the user's job run during the backup.

Temporarily Suspending Jobs

It is advisable to suspend the execution of all jobs while a STORE is being performed. (If a job step attempts to open a file for write access which is being stored, the open fails, with possible negative consequences for the job.)

NOTE Advise users not to schedule jobs during the time schedule for backups. STORE does not copy files opened by suspended jobs. Therefore, a backup undertaken with suspended jobs may be incomplete.

To check whether there are any jobs running on the system, enter:

```
: SHOWJOB JOB=@J
```

If you see a message stating NO SUCH JOBS, or if the list contains no executing jobs (check the STATE column for EXEC), skip to the "Send Another Warning" step below. If any jobs are in the EXEC state, you should temporarily stop them. To do so, use the job numbers that appear in the first column on your screen with the BREAKJOB command, as follows:

```
: BREAKJOB #Jnnn
```

Repeat the BREAKJOB command for each executing job in the list. When you have suspended all executing jobs, check their status by entering:

```
: SHOWJOB JOB=@J
```

Each job should now be listed as SUSP, or suspended. However, jobs that have resources allocated to them do not suspend immediately; you may have to wait for them to suspend.

NOTE For an online backup, suspend all active jobs until you receive the message on the console telling you that files can now be modified. At that point, you can resume any jobs that you suspended before the backup starts.

For a 7x24 true-online backup, you do not need to suspend any active jobs before starting the backup. They can continue to run throughout the backup.

Sending Final Warning Messages

Check to see whether any users are still logged on to the system by entering:

```
: SHOWJOB JOB=@S
```

If there are any remaining sessions, besides your own, send them one last warning:

```
: WARN @S; LOG OFF NOW! BACKUP ABOUT TO BEGIN
```

NOTE For an online backup, issue this warning to all users who are still logged on. You may want to add that they can log back on once the backup has started. As noted above, the amount of time they need to remain logged off depends on the amount of data being stored.

For a 7x24 true-online backup, you can send a warning message to all users. However, you do not need to tell them to log off. It may be a good idea to warn users of an impending 7x24 true-online backup, since they may see some decrease in system response time, due to system resources being used for the backup.

Aborting Any Remaining Sessions

Check one more time whether all sessions (except your own) have logged off by entering:

```
:SHOWJOB JOB=@S
```

Either abort any sessions other than your own that are still executing, or check with the users and ask them once more to log off. To abort sessions, use the `ABORTJOB` command and the session numbers that appear in the first column on your screen, as follows:

```
:ABORTJOB #Snnn
```

After you have aborted all sessions (except your own) issue a `SHOWJOB` command to check that all sessions are aborted. When the system responds that you are the only session on the system (as shown below), you are ready to begin the backup.

Figure 4-1. Aborted Session System Response

```
JOBNUM  STATE  IPRI  JIN  JLIST  INTRODUCED  JOB NAME
#S184   EXEC           20  20    FRI 1:03P  OPERATOR.SYS
#J17    SUSP           10S  PP    FRI 5:02P  AJOB,MRS.T
2 JOBS:
  0 INTRO; 0 SCHEDULED
  0 WAIT;  INCL 0 DEFERRED
  1 EXEC;  INCL 1 SESSIONS
  1 SUSP
JOBFENCE= 14; JLIMIT= 0; SLIMIT= 0
```

NOTE For an online backup, abort any remaining sessions on your system before starting the backup. You can allow jobs and users back on the system after receiving the console message from TurboSTORE indicating that files can now be accessed.

For a 7x24 true-online backup, you do not have to abort any remaining jobs or sessions. They can continue to run throughout the backup.

5 Storing Files

This chapter contains complete information on how to use the `STORE` command to store files to backup media. `STORE` can be used to back up a single file, all user, system, and system directory files (that is, a full system backup), or any number of files in between.

The information in this chapter describes how to use the `STORE` command and its options to create backups that can be restored using the `RESTORE` command. The subjects are:

1. Specifying files in the `STORE` command
2. Using indirect files
3. Using filesets with the `RENAME` option
4. Storing database files
5. Storing to labeled tapes
6. Monitoring `STORE` command progress

Use the `STORE` command, and optionally the `FILE` command, to store a file or set of files. A `FILE` command is used to assign a filename to a backup device. A `STORE` command names the files to store and the options to use.

The `STORE` command options for naming files let you name a single file, several files, or all the files on the system. The group of files to store is called a fileset. In addition to naming many kinds of filesets, you can name filesets to exclude from the backup.

Refer to Chapter 6, "STORE and TurboSTORE Options," for a description of each of the `STORE` options. For example, you can print a list of all files stored, purge disk files immediately after storing them to backup, or request progress messages that show the status of the `STORE` process on the console.

Additionally, the `MPE/iX` command, `VSTORE`, allows you to verify that the data written to backup using the `STORE` program is valid. Refer to Chapter 9, "When STORE is Complete," for a description of using `VSTORE`, and Chapter 11, "Restoring Files," for descriptions of the `VSTORE` options, which are essentially the same as the `RESTORE` options.

Files can be stored to a single backup device or to multiple backup devices in several configurations. You can also store files in interleave format. Use multiple backup devices and interleave format to increase performance when you have a large set of files to store, for example, when you are performing a full system backup.

See Chapter 3 for information on preparing backup devices.

CAUTION Do not use the `STORE` or `RESTORE` programs to transfer any of the files that originally came from your system load tape (SLT) from LDEV 1 to a disk

other than LDEV 1. These files reside in the SYS account and can be displayed by entering the SHOW command in the SY configurator of the SYSGEN utility.

Also, do not purge or put a lockword on these files. Finally, do not use any command other than SYSGEN to manage the files within a configuration group. These configuration files come as a set and individual files may not be moved between groups or systems. Such actions will make these system files inaccessible causing problems either immediately or at a later time.

Specifying Files in a STORE Command

The files that you are allowed to store depend on your capability. In general, you may only store files that you have read access to. If you are a system manager, system supervisor, or system operator, you can store any file in the system.

The following section describes how you specify which files you want to be stored. The entire set of files to be stored is known as a *fileset list*.

The fileset list is the first argument to the store command. If no files are specified, the default set is all files in your current working directory, or "@". This is the default no matter what capabilities you may have.

-
1. If an empty fileset list is specified, and the DIRECTORY option is used, the default set is empty. This allows a system administrator to create a backup containing just the system accounting structures, without storing any files.
 2. If a fileset list is not specified, a semi-colon (;) place holder must be used instead. For example:

```
:FILE T;DEV=TAPE
:STORE ;*T
```

If a fileset list is specified, it must follow certain guidelines. The list is made up of comma separated fileset items. Each item can be either a fileset or an indirect file.

Filesets

A fileset specifies a group of files to be stored. It can also include an exclusion set, which is a group of files to be excluded from the STORE operation.

The syntax for a fileset is:

```
filestore[- filestoexclude [- .... ] ]
```

Where filestore represents a fileset and filestoexclude represents an exclusion fileset.

The filestore and filestoexclude are specified using the same syntax. Any file that

matches `filestore` will be stored, unless the file also matches a `filestoexclude`. An unlimited number of `filestoexclude` may be specified. However, if the `TRANSPORT` option is also being used, only one `filestoexclude` may be specified.

NOTE Since the hyphen (-) is a valid character for HFS syntax filenames, a blank character must separate it from HFS filesets to obtain the special negative fileset meaning.

Specifying A Fileset

Filesets and exclusion sets are specified using the same format. They can be specified two different ways:

- MPE syntax
- HFS syntax

MPE Syntax

A file is specified as:

```
filename[.groupname[.accountname] ]
```

If `accountname` is omitted, the file is looked for in your logon account. If `groupname` is omitted, the file is looked for in your current working directory.

A lockword may be specified for a file, in the form:

```
filename/lockword.group.account
```

HFS Syntax

A file is specified as:

```
/dir_lev_1/dir_lev_2/.../dir_lev_i/.../filedesig
```

or

```
./dir_lev_i/dir_lev_j/.../dir_lev_k/.../filedesig
```

If the name begins with a dot (`.`), then it is fully qualified by replacing the dot with the current working directory (CWD).

Each of the components `dir_lev_i` and `filedesig` can have a maximum of 255 characters with the full pathname being restricted to 1023 characters.

Each of the components `dir_lev_i` and `filedesig` can use the following characters:

- Letters a to z
- Letters A to Z
- Digits 0 to 9
- Special characters - _ .

For HFS syntax, the lowercase letters are distinct from the uppercase letters (no upshifting). Names in MPE syntax are upshifted.

Using Wildcard Characters

Use wildcard characters to select multiple files to store. The wildcard character is expanded to include all files that match.

Both MPE and HFS name components use the at sign (@), pound sign (#), and question mark (?) as wildcard characters. These wildcard characters have the following meaning:

@	Specifies zero or more alphanumeric characters.
#	Specifies one numeric character.
?	Specifies one alphanumeric character.

The characters can be used as follows:

<i>n</i> @	All files starting with the character <i>n</i> .
@ <i>n</i>	All files ending with the character <i>n</i> .
<i>n</i> #####	All files starting with character <i>n</i> followed by seven digits (useful for storing all EDIT/3000 temporary files).
<i>n</i> @ <i>x</i>	All files starting with the character <i>n</i> and ending with the character <i>x</i> .
? <i>n</i> @	All files whose second character is <i>n</i> .
<i>n</i> ?	All two-character files starting with the character <i>n</i> .
? <i>n</i>	All two-character files ending with the character <i>n</i> .

Also, character sets may be specified in the following syntax:

[<i>ct</i>]	Specifies letter <i>c</i> or <i>t</i> .
[<i>c-t</i>]	Specifies any letter from range <i>c</i> to <i>t</i> .
[<i>e-g1</i>]	Specifies any letter range <i>e</i> to <i>g</i> or digit <i>1</i> .

Examples of using character sets are:

[<i>A-C</i>]@	All files that begin with the letters <i>A</i> , <i>B</i> , or <i>C</i> .
myset[<i>e-g1</i>]	All files that begin with the name <i>myset</i> and end in <i>E</i> , <i>F</i> , or <i>G</i> , or <i>1</i> .
myset[<i>d-e1-6</i>]	All files that begin with the name <i>myset</i> and end in <i>D</i> or <i>E</i> , or <i>1</i> , <i>2</i> , <i>3</i> , <i>4</i> , <i>5</i> , or <i>6</i> .

You specify up to a maximum of sixteen characters for each character set, and you cannot nest brackets. Do not use character sets with the `TRANSPORT` option.

A character set specifies a range for only one ASCII character. The range [*a-d*]@ gets all files that begin with the letter "a" through the letter "d". The range [*ad-de*] will cause unpredictable results.

NOTE Since the hyphen (-) is a valid character for HFS syntax filenames, it is allowed inside a character set, immediately following a left bracket ([) or preceding a right bracket (]). When specified between two characters, the hyphen implies a range of characters.

MPE and HFS Naming Equivalences

When an MPE name component is a single @ wildcard, the @ includes all MPE and MPE-named files at that level and below. To specify only MPE-named files, use ?@ instead. The ? forces the first character of the filename to be an alphanumeric character, which only matches MPE-named files.

MPE wildcards are not expanded in exclusion filesets. This means that @.@.@-@.@.@ is NOT an empty fileset. @.@.@ is the same as /. -@.@.@ does not have an HFS equivalent. -@.@.@ specifies all HFS-named files on the system, the resulting fileset is all of the HFS named files on the system.

You can enter a fileset in any of the following formats and may use wildcard characters. Equivalent MPE and HFS formats are grouped together.

file.group.acct	One particular file in one particular group in one particular account.
file.group	One particular file in one particular group in the logon account.
file ./file	One particular file in the CWD.
@.group.acct /ACCT/GROUP/	All files (MPE and HFS) in one particular group in one particular account.
?@.group.acct	All MPE-named files in one particular group in one particular account.
@.group /LOGONACCT/GROUP/	All the files (MPE and HFS) in one particular group in the logon account.
?@.group	All MPE-named files in one particular group in the logon account.
@.@.acct /ACCT/	All the files (MPE and HFS) in all the groups in one particular account, plus all the files and directories under the specified account.
?@.@.acct	All MPE-named files in all the groups in one particular account.
thisisit.@.acct	Any MPE-file named thisisit in all the groups in one particular account.
@	All (MPE and HFS) files in the CWD. This is the default for everyone, regardless of permissions.
@.@	All (MPE and HFS) files in the logon account.
?@.@.@	All MPE-named files in the system.
@.@.@ /	All the files and directories (MPE and HFS) in the system.

Examples:

NOTE All of the following examples assume that a file equation for T was previously set up, by issuing the FILE command:

: FILE T ; DEV=TAPE

- Storing a single file

: STORE FILE1.PUB.RESEARCH;*T

- Storing several files
 - `:STORE FILE1.PUB.RESEARCH,FILE2.PUB.MYACCT;*T`
- Store several files with wildcards
 - Store all MPE and HFS files in GROUP1.ACCTNAME:
 - `:STORE @.GROUP1.ACCTNAME;*T`
 - Store all MPE and HFS files in the dir3 directory:
 - `:STORE ./dir3/@;*T`
 - Store all HFS files that start with lower case letters in the CWD:
 - `:STORE ./[a-z]@;*T`
- Exclude some files from a fileset:
 - Store everything on the system except files in PUB.SYS, MPE format:
 - `:STORE @.@.-@.PUB.SYS;*T`
 - HFS format:
 - `:STORE / - /SYS/PUB/@;*T`
- Store everything on the system except the 3000devs account:
 - `:STORE @.@ - /3000devs;*T`
- Store only the HFS-named files on the system:
 - `:STORE / - @.@.*;*T`
- Store only the MPE-named files on the system:
 - `:STORE ?@.@.*;*T`
- All files in the TEST group, except those that begin with A or B:
 - `:STORE @.TEST - A@.TEST - B@.TEST;*T`

Using Indirect Files

If you back up the same information regularly, you might keep the STORE parameters and options you use in an indirect file. An indirect file is a text file containing filesets, exclusion sets, and parameters for a STORE command. Instead of listing the files you want to store and the options you want to use in a STORE command, enter them in a text file and name the text file in the STORE command.

An indirect file must:

- Be an existing permanent or temporary file.
- Have a record length between 8 and 14K bytes.
- Have read access allowed, although the file may be open shared or intrajob.

The following is an example of using an indirect file:

Suppose you regularly back up all the files on your system, including the system directory

and a nonsystem volume set's files and directory. The usual STORE command, as illustrated in the previous section, might be:

```
: STORE @. @. @; *T; DIRECTORY; ONVS=MPEXL_SYSTEM_VOLUME_SET, NV1
```

Using an indirect file can save time, especially if your STORE command is long or contains information that is difficult to remember. You enter the STORE files and parameters in a text file:

```
@. @. @; DIRECTORY; ONVS=MPEXL_SYSTEM_VOLUME_SET, NV1
```

Give the text file a name that is easy to remember, and reference the text file name in a STORE command. The following STORE command references the file BACKUP as an indirect file:

```
: STORE ^BACKUP; *T
```

The caret (^) is used prior to the filename to tell STORE that it is an indirect file.

The exclamation point (!) can also be used, but it will be interpreted by the command interpreter (CI) as a variable deference. If you do not intend for it to be used that way, you should specify two exclamation points (!!) to cause a single exclamation point (!) to be passed to STORE. However, the recommended character is the carat (^).

Multiple indirect files can also be used. For example:

```
: STORE ^BACKUP1, ^BACKUP2; *T
```

Indirect files can also be nested. The maximum nesting level is 3. For example:

```
INDIR1: @. PUB. SYS; SHOW
```

```
INDIR2:
@. OUT. HPSPPOOL
^INDIR1
```

```
: STORE ^INDIR2; *T
```

Any commands that can be specified on the STORE command line can be specified in an indirect file. The one exception is the store device specification, which is usually specified immediately following the fileset list. The store device must always be specified on the command line. However, if multiple devices are being used for the backup, the STORESET option can be specified in the indirect file.

Options can be specified on both the command line and in the indirect file. Any options that are specified apply to all filesets in all indirect files, as well as any filesets specified on the command line.

Each line in the indirect file can contain fileset lists separated by commas, and options separated by semicolons. Fileset lists must always be specified first on a line. An option can also be specified first on a line, but once an option is specified, no fileset lists may follow them on the line.

Negative filesets can also be specified in indirect files. They should always immediately follow a fileset specification. Although one or more negative filesets can be specified on their own line, no options should be specified between the fileset and its associated negative filesets. As when specifying negative filesets on the command line, there is no

limit to the number of negative filesets that can be specified.

Do not specify any continuation characters, such as ",", ";", "+", or "&", at the end of an indirect file line. If you need to specify more file sets than can fit on a line, just continue specifying them on additional lines. Do not start a line with the "," character. CI variables are not expanded in indirect files, so do not specify filesets using them.

Blank lines are **not** allowed in indirect files.

An example of a more complicated indirect file is:

```
@.MYGR.MYACCT
- A@.MYGR.MYACCT
- /MYACCT/MYGR/B@;SHOW
@.MYGR2.MYACCT - A##.MYGR2.MYACCT
;PROGRESS
@.MYGR3.MYACCT, @.MYGR4.MYACCT
```

Using Filesets With the RENAME Option

If the `RENAME` option is specified, the fileset syntax is expanded to allow new filenames and creators to be specified. The `RENAME` syntax is:

```
filestore [- filestoexclude [- ...] ][=targetname]
```

The target name field specifies the new name and creator for the file on the store media. It has the form:

```
filename [:creator[.creatoraccount]]
```

The filename can be any legal MPE filename or HFS pathname. The creator and creatoraccount must be legal creator and account names, respectively. The only wildcard character allowed is a single at sign (@) for each component of the filename, creator or creatoraccount. The wildcard character @ indicates that the source value for that component should be used. An HFS pathname which ends in a / is considered an HFS directory and no wildcard characters are allowed.

The `RENAME` option must be specified if the targetname is used. Refer to the `RENAME` section in Chapter 6, "STORE and TurboSTORE Command Options," for more details.

Storing Symbolic Links and Device Links

As of Release 5.0 and later, MPE/iX supports symbolic links and device links. When storing symbolic links, TurboSTORE only follows the link if it is not the last component of an HFS filename. If the symbolic link is the last component of an HFS filename, the only the link, or "container" will be stored. When storing device links, only the link is stored.

For example, if you have a symbolic link named `/SYS/TESTDIR` which points to the directory `/MYACCT/MYDIR`, then issuing the following `STORE` command would store all of the files in `/MYACCT/MYDIR`:

```
STORE /SYS/TESTDIR/@; ;SHOW
```

If you had a symbolic link name `/SYS/PUB/MYFILE`, which pointed to `/MYACCT/PUB/TARGET`, then the following `STORE` command would only store the symbolic link `MYFILE`, not the target file `TARGET`:

```
STORE /SYS/PUB/MYFILE; ;SHOW
```

For more information regarding symbolic links and device links, see the article "Symbolic Links on MPE/iX" in Chapter 5 of the *Communicator 3000 MPE/iX General Release 5.0 (Core Software Release C.50.00)*.

Storing Database Files

Storing TurboIMAGE Databases

Prior to MPE/iX Release 5.5, the only way to logically back up databases was by using the `DBSTORE` and `DBRESTOR` commands. However, the `DBSTORE` program only allowed certain `STORE` options, and could only store or restore one database at a time. The features of `DBSTORE` are now available when using `STORE`, `TurboSTORE/iX II` or `TurboSTORE 7x24 True-Online Backup`.

To specify a complete database to be stored, only the root filename needs to be specified. (By default, `STORE` does **not** allow you to specify an incomplete TurboIMAGE database to be stored when you use the 7x24 true-online backup option `ONLINE=START` or `ONLINE=END`.) When `STORE` encounters a database root file, it will use the database to determine what other files are part of the database. Third-party indexing files will also be included in this list, if available.

If a wildcard that includes an entire database is specified, `TurboSTORE/iX` will still verify the file list with the database to ensure that all database files are being written to the backup.

If database dataset files are specified for a backup, but their root file is not specified, those dataset files will not be stored. For each file specified, the following warning will be displayed:

```
MUSIC03.MUSIC.MYDB NOT STORED: FILE IS PART OF AN IMAGE  
DATABASE AND ROOT IS NOT SPECIFIED.
```

If any of the dataset files are specified in addition to the root file, no warning will be displayed but the individual dataset files will be counted as redundantly specified files.

If you are using the `TurboSTORE/iX 7x24 true-online backup ONLINE` option, then all TurboIMAGE databases fully specified, either by wildcards or by specifying the root file, are quiesced at the online backup sync point. If you want to back up part of an IMAGE database, you can specify the `PARTIALDB` option. For more information on `PARTIALDB` and storing databases with 7x24 true-online backup, see "Storing Database Files Using `PARTIALDB` or `FULLDB`" in Chapter 6, "STORE and TurboSTORE Command Options."

CAUTION If some of the dataset files are missing from the database, the quiesce may not succeed. `STORE` will still back up the dataset files listed in the fileset list.

Storing ALLBASE/SQL Databases

A similar convention exists for ALLBASE/SQL databases. Only the DBE file for an ALLBASE/SQL database needs to be specified. `STORE` will query the database to get the complete list of files to be stored.

If an individual ALLBASE/SQL database file is specified without the DBE file, it will not be stored, and the following message will be issued:

```
MYSET01.MYDB.SYS NOT STORED: FILE IS PART OF AN ALLBASE DB ENVIRONMENT  
AND DBCON IS NOT SPECIFIED.
```

If any of the dataset files are specified in addition to the root file, no warning will be displayed but the individual dataset files will be counted as redundantly specified files.

If you are using the TurboSTORE/iX 7x24 true-online backup `ONLINE` option, then all ALLBASE/SQL databases fully specified, either by wildcards or by specifying the root file, are quiesced at the online backup sync point. If you want to back up part of an ALLBASE DBEnvironment database, you can specify the `PARTIALDB` option. For more information on `PARTIALDB` and storing databases with 7x24 true-online backup, see "Storing Database Files Using `PARTIALDB` or `FULLDB`" in Chapter 6, "STORE and TurboSTORE Command Options."

CAUTION If some of the dataset files are missing from the database, the quiesce may not succeed. STORE will still back up the dataset files listed in the fileset list.

Storing Non-HP Databases

When storing databases other than than TurboIMAGE/SQL and ALLBASE/SQL databases, STORE does not enforce any root file/data set rules. You should continue to store other databases using the same procedures used prior to MPE/iX Release 5.5.

Storing to Labeled Tapes

The `STORE` command syntax for backups to labeled tapes is identical to that for unlabeled tapes. However, you must use a `FILE` equation to request that the `STORE` command write to a labeled volume. For example, when preparing for a backup to a volume set labeled `VOL25`, to be mounted on logical device 7, the following file equation could be entered:

```
:FILE TAPE7=BACKUP.JAN29;DEV=TAPE;LABEL="VOL25",ANS,02/28/90
```

This file equation is followed by a `STORE` command which references the destination device:

```
:STORE @.@@;*TAPE7;storeoptions
```

Using Previously Labeled Volumes

If the previously labeled volume set has not been mounted prior to the backup, the system sends you a message requesting the first volume of the set. For example:

```
?23:20/#S5/50/Mount volume VOL25 (ANS)
```

After you mount the correct volume and it has been recognized by automatic volume recognition (AVR), you will see a message similar to the following:

```
23:23/15/Vol VOL25 (ANS) mounted on LDEV# 7
```

After this volume has been successfully written to, `STORE` requests that the next reel of the set be mounted. For example:

```
?23:34/#S5/50/MOUNT REEL 2 OF VOLSET VOL25 ON LDEV 7
```

NOTE AVR messages for all reels beyond the first reel are issued by `STORE`. These messages cannot be recalled.

Using Unlabeled Volumes

If you plan to create labeled volumes using backups that are *not* labeled, the preceding procedure differs slightly. Use the `FILE` equation and `STORE` command as you would with a labeled volume. For example:

```
:FILE TAPE7=BACKUP.JAN29;DEV=TAPE;LABEL="VOL25",ANS,02/28/90
```

```
:STORE @.@@;*TAPE7;storeoptions
```

You will receive a message similar to the following:

```
?23:26/#S/47/MOUNT VOLUME VOL25 (ANS)
```

This is followed by a message recognizing the media that is mounted:

```
23:27/15/Vol (unlabeled) mounted on LDEV# 7
```

Reply to the first message by indicating the `LDEV` where the volume is to be mounted. For example:

```
:=REPLY 47,7
```

The volume mounted is then labeled and the backup continues until a new volume is required. You will receive a message similar to the following:

```
?23:33/#S5/47/MOUNT REEL 2 OF SET VOL25 ON LDEV 7
```

This message is followed by a request to provide a volume ID for the volume.

```
?23:34/#S5/47/REPLY WITH VOL ID FOR VOLSET VOL25 on LDEV 7?
```

Reply to this message with a valid ID for the volume. For example:

```
:=REPLY 47,VOL26
```

The backup continues, requesting new volumes until all files in the fileset are written to the backups.

Monitoring STORE Command Progress

After you enter a STORE command, STORE issues a message similar to the following one on the \$STDLIST:

```
TURBO-STORE/RESTORE VERSION A.20.01 (C) 1986 HEWLETT-PACKARD CO.  
TUE, NOV 6, 1989, 2:57 PM
```

Monitor the console for additional messages. In addition to regular progress messages, if you requested them with the PROGRESS option, STORE sends you messages when you need to reply to a backup request, add a write ring to a backup reel, or mount a piece of media on a backup device.

Responding to Tape Requests

When your backup devices are not configured to be automatically allocated, the system answers your STORE command with one or more backup requests. Refer to Chapter 8, "Handling Media Requests," for more information.

If you have previously issued a FILE command that included the backup device's LDEV number, the backup request asks you to confirm that number. For example:

```
?14:57/#S25/43/IS "T" ON LDEV #8 (Y/N)?
```

Replying to Tape Requests for a Multiple-Device Backup

If you are backing up files to several devices, the system sends you a backup request for each device. If you give your backup devices unique filenames, you receive a backup request for each unique filename. Use the "Multiple Device Planning Chart" from figure 3-7 in Chapter 3, "Preparing Storage Devices," to determine which LDEV to assign to each name.

If you give your backup devices a common filename, you receive several identical backup requests referring to the common filename, one for each device. For example, the first backup request in a sequential backup asks for the LDEV number of the first device. The second request asks for the LDEV number of the second device, and so on.

During a parallel backup, backup requests for the parallel devices will be issued to the console. Requests for all parallel devices will be issued at once, in no specific order. There will be a different PIN associated with each request, although all requests are really coming from STORE. You should reply to the requests using the appropriate LDEVs for the parallel sets. Files will not begin to be stored until all requests have been replied to.

When you are using parallel device pools, the requests for the first serial device in all pools will be issued first, in no specific order. After the first serial device has been replied to, then a request will be issued for the next serial device, and so on. Parallel pool requests are independent, with each parallel pool having a different PIN associated with its request. Files will not be stored until all requests have been replied to.

Handling Media Errors

When a media error occurs while writing to an output device, the behavior of STORE depends on the value of the `ONERROR` option, as explained below.

Tape Error Recovery

The STORE command's `ONERROR` parameter lets you specify a backup error recovery procedure. Your options are `ONERROR=QUIT` and `ONERROR=REDO`. `REDO` is the default. `QUIT` instructs STORE to abort upon encountering a backup I/O error while `REDO` instructs STORE to continue after encountering an I/O error.

If you specify `REDO` and the STORE command encounters a backup I/O error on any device, STORE sends a message to your terminal (or the file or device to which you have assigned the file `SYSLIST`). This message indicates where the error occurred. For example, `STORESET=(*T1 , *T2 , *T3)` is requested and a backup error occurs on the device referenced by `*T2` while it is creating reel 2. STORE rewinds the reel to the load point, marks the reel as bad, and asks you to mount another reel. After you mount the new reel, STORE automatically stores all the files from the point where the bad reel began.

The message is similar to the following:

- The user sees:

```
STORE ENCOUNTERED MEDIA WRITE ERROR ON LDEV #
```

where `LDEV #` is the device referenced by

- `T2`. This means that the store error recovery occurs on the same device where the error occurred.
- The operator sees the following messages on the console:

```
STORE IS MARKING REEL BAD ON LDEV #
```

```
MOUNT MEDIA # OF SET # FOR STORE ON LDEV #
```

Bad Media

Once a piece of media has been marked bad, it should not be used for future STORE operations. If a media that has been marked bad is mounted for use in a STORE operation, STORE will issue the following message to the console:

```
MEDIA ON LDEV # IS BAD. CONTINUE ANYWAY?(Y/N)
```

If you reply `YES`, then STORE will write over this media. If you reply `NO`, the media will be unloaded, and you should mount another piece of media.

NOTE STORE will not check to see if a piece of media has been marked bad if it is the very first reel mounted. Therefore, it is important that once a reel has been marked bad by store, some note is made on the physical label of the media to indicate that it encountered a write error while storing to it.

Responding to a NO WRITE RING Message

Reel backups require you to insert a write ring into the reel in order to store files onto the backup. You see a NO WRITE RING message when you attempt to store files to a reel-to-reel backup and neglect to insert a write ring into the groove on the back of the backup reel. For example:

```
15:57/2/LDEV#8 NO WRITE RING
```

The message above tells you there is no write ring on the backup reel mounted on LDEV 8. Without the write ring, the files on the backup are protected from being inadvertently copied over. You must rewind and remove the backup reel, and mount a backup which has a write ring. Follow the instructions in Chapter 2, "Managing Backups and Backup Media." When the backup is correctly mounted, STORE continues, and another REPLY is not necessary.

You can also get this message when attempting to store to a write-protected DDS tape or 3480 cartridge. Follow the instructions in your device's manual for making your media writeable. Usually this involves sliding a tab on the end of the tape (for DDS), or turning a dial on the end of the cartridge (for 3480).

Progress Messages

If you have used the PROGRESS parameter in your STORE command, the system displays progress messages at the interval you specified. For example:

```
STORE OPERATION IS 4% COMPLETE
```

You see similar messages at regular intervals, for example:

```
STORE OPERATION IS 7% COMPLETE
STORE OPERATION IS 11% COMPLETE
STORE OPERATION IS 14% COMPLETE
```

The backup is finished when the 100% complete message displays.

Mounting Additional Tapes

When STORE fills a piece of media on a device and it has additional information to store to that device, it asks you to mount another piece of media. You see a mount request and LDEV NOT READY message like the ones below:

```
15:28/#S43/59/MOUNT MEDIA FOR STORE ON LDEV 8
15:29/#S43/59 LDEV 8 NOT READY
```

The mount request above asks you to mount another piece of media on logical device 8.

To mount another piece of media:

1. Unload the first piece of media and stack it, label side down, on top of any media previously written on this drive during the current STORE operation. Remove the write ring.
2. Select the next piece of media from the prepared stack near the piece of media device. If there are no more media prepared, prepare a new piece of media following the instructions in "Preparing Your Backup Media", in Chapter 3, "Preparing Storage

Devices."

3. Mount the new piece of media.
4. Continue monitoring the backup process. Mount new media as the system requests them.

6 STORE and TurboSTORE/iX Command Options

This chapter describes how to use all of the options available with all STORE and TurboSTORE/iX products, as of MPE/iX Release 5.5. The following options are new as of MPE/iX Release 5.5:

```
STATISTICS
PARTIALDB
FULLDB
STOREDIR[ ECTORY ]
NOSTOREDIRE[ ECTORY ]
LOGVOLSET
```

All the STORE and TurboSTORE/iX product options are listed with brief descriptions below. They are in the same order displayed in the STORE command syntax in Appendix B. The options discussed are:

STORE Options

Option	Description
SHOW	Lists the files stored
ONERROR	Specifies error recovery method
DATE	Stores files specified by a date: files modified since a specified date or files not accessed since a specified date
PURGE	Removes stored files from the system for archiving
PROGRESS	Displays progress messages during backup
DIRECTORY	Stores a backup of the system directory
FILES	Specifies the number of files to store
TRANSPORT	Creates compatible backups in two formats: MPE V/E format and pre-POSIX format
COPYACD, NOACD	Specifies file security
FCRANGE	Stores files with specified file codes
MAXTAPEBUF	Improves storing efficiency by increasing the record size of the backup device
NOTIFY	Allows STORE to stream a user specified job file
STATISTICS	Provides extra data about a backup

ONVS	Stores files from volume sets
SPLITVS	Stores files located on mirrored volume sets
RENAME	Changes the name of original files on the system or the target files on the backup media
TREE, NOTREE	Specifies if files should be scanned recursively to include all files in a hierarchical directory
INTER	Specifies more than one file at a time should be written, or <i>interleave</i> on the backup media.
STORESET	Allows you to specify multiple backup devices.
PARTIALDB	In a 7x24 true-online backup environment, allows you to back up part of a database fileset. Default for a non-7x24 true-online backup (that is, <code>ONLINE=START</code> or <code>ONLINE=END</code> is not specified).
FULLDB	In a non-7x24 true-online backup environment, allows you to backup the entire database by just specifying the root filename only in the fileset. Default for a 7x24 true-online backup.
STOREDIRECTORY	Allows you to create a disk copy of the STORE directory of the backup. This is enabled by default when <code>ONLINE=END</code> is specified. The disk directory allows RESTORE to more quickly locate files on the backup.
NOSTOREDIRE[CTORY]	Allows you to prevent creating a STORE disk directory file.

TurboSTORE/iX II Options

`COMPRESS` TurboSTORE option that improves storing efficiency, compresses data being stored

TurboSTORE/iX 7x24 True-Online Backup Options

`ONLINE` Option that stores files while they are open.

`LOGVOLSET` Option that specifies which volume set to use for log files.

Note that all options available in FOS are also available in TurboSTORE/iX II and TurboSTORE/iX 7x24 True-Online Backup. All options available in TurboSTORE/iX II are also available in TurboSTORE/iX 7x24 True-Online Backup.

- TurboSTORE/iX II provides you with certain special features not available with the standard STORE program.
- TurboSTORE/iX 7x24 True-Online Backup enables you to perform online backups. This allows you to back up files and applications without closing files or disrupting users.

This chapter provides detailed information on the options for all STORE and TurboSTORE products.

STORE Options

Listing the Files Stored (SHOW)

Whenever you store a set of files, the system displays the total number of files stored at your terminal. If there was an error, it also displays the number and names of files not stored and the reason each was not stored. Use the `SHOW` parameter and its options to display the names and additional information about the files stored and to list them on your system printer as well as at your terminal.

Internally, the `STORE` program writes information about the files it stores to a file with the formal file designator `SYSLIST`. It prints file information at your terminal by equating `SYSLIST` with the standard listing device for your session (your terminal). You can redirect the file information to another file or device by issuing a `FILE` command that assigns `SYSLIST` to that device or file. For example, you might want to redirect the file information to a disk file so that you can keep it online for your records.

To write `STORE` information to the disk file, `SHOWFILE`, enter:

```
:FILE SYSLIST=SHOWFILE;DEV=DISC
```

To print information on the files you stored at the system printer as well as at your terminal, use the `SHOW=OFFLINE` option:

```
:FILE T;DEV=TAPE  
:STORE @.OPERATOR.SYS;*T;SHOW=OFFLINE
```

To redirect the `OFFLINE` listing to another file or device, issue the following `FILE` command, it will send a copy of the `STORE` listing to the disk file `SHOWOFFL`:

```
:FILE OFFLINE=SHOWOFFL;DEV=DISC
```

To generate a store listing to the terminal as well as to a disk file, enter the previous `FILE` command followed by the `STORE` command:

```
:STORE @.OPERATOR.SYS;*t;SHOW=OFFLINE
```

Using a Long or Short File Display

The `STORE` command gives you two options for displaying additional information about the files stored. Using the `SHOW` parameter, you can choose to display file information in either long or short form.

A short file display prints the file's full name, volume restrictions, number of sectors, file code, and media number for each file you store.

In addition to there being a long and a short `SHOW` listing, there are also two different versions of each - MPE and HFS. The type of listing printed depends on the input fileset list. You will get a HFS style listing if:

- At least one file in the fileset list is specified in HFS syntax.
- A fileset wildcard is expanded to include a HFS syntax file.

The HFS style listing contains the same information as the MPE style listing. The only difference is that the filename is printed as an absolute pathname at the end of the line.

Since HFS pathnames can be variable lengths, and can be quite long, they are placed at the end of the line so that they do not affect the spacing of the other columns. If the name does not fit on one line, it wraps to the next line. An asterisk (*) is placed at the last column of the line to indicate that the rest of the filename continues on the next line.

NOTE It is likely that a STORE with the fileset @.@.@ will result in a HFS format listing, since MPE/iX, by default, contains some HFS-named files.

To display the list of files stored using the short form in MPE format, use the SHOW=SHORT option in your STORE command. For example:

```
:FILE T;DEV=TAPE
:STORE @.PUB.TEST;*T;SHOW=SHORT
```

Figure 6-1. illustrates a short file display in MPE format using the above example.

Figure 6-1. Short File Display (MPE Format)

FILENAME	GROUP	ACCOUNT	VOLUME	RESTRICTIONS	SECTORS	CODE	MEDIA
FILE1	.PUB	.TEST	DISC	:C	1104		1
FILE2	.PUB	.TEST	DISC	:C	0		1
FILE3	.PUB	.TEST	DISC	:C	32		1

This example uses the same fileset as the first example, except that it is specified in HFS syntax.

```
:FILE T;DEV=TAPE
:STORE /TEST/PUB/@;*T;SHOW=SHORT
```

Figure 6-2. illustrates a short file display in HFS format using the above example.

Figure 6-2. Short File Display (HFS Format)

VOLUME	RESTRICTIONS	SECTORS	CODE	MEDIA	FILENAME
DISC	:C	1104	1		/TEST/PUB/FILE1
DISC	:C	0	1		/TEST/PUB/FILE2
DISC	:C	32	1		/TEST PUB/FILE3

This example shows how the HFS syntax names will be wrapped, if they cause the line to exceed 80 characters:

```
:FILE T;DEV=TAPE
:STORE /usr/lib/terminfo/e/e[t-z]@;*T;SHOW=SHORT
```

Figure 6-3. illustrates a short file display in HFS format with line wrap using the above example.

Figure 6-3. Short File Display with Wrapped HFS Names

VOLUME	RESTRICTIONS	SECTORS	CODE	MEDIA	FILENAME
DISC	:C	256	1		/_HFSACCT/_HFSGRP/HFSMAP
DISC	:C	16	1		/usr/lib/terminfo/e/ethern*

```

et
DISC          :C      16      1      /usr/lib/terminfo/e/ex3000
DISC          :C      16      1      /usr/lib/terminfo/e/exidy
DISC          :C      16      1      /usr/lib/terminfo/e/exidy2*
500

```

A long file display contains the same information as a short file display plus each file's ending media number, record size, blocking factor, maximum number of extents allowed, end-of-file, and file record limit.

Once again, there are two formats for the long listing -- MPE and HFS. The HFS format contains the same information as the MPE format, except that the variable-length HFS filename is at the end of the line. It will be wrapped in the same way as the short listing if the filename length exceeds the line length.

To display the list of files stored using the long form in MPE format, use the `SHOW=LONG` option. For example:

```

:FILE T;DEV=TAPE
:STORE @.PUB.TEST;*T;SHOW=LONG

```

Figure 6-4. illustrates a long file display from the above example.

Figure 6-4. Long File Display (MPE Format)

```

      FILENAME GROUP ACCOUNT VOLUME RESTRICTIONS SECTORS CODE MEDIA_NUM RSIZE BFCTR
#EXT EOF    LIMIT
      FILE1  .PUB  .TEST  DISC      :C      1104      1 - 1      80   16   0
3500 3500
      FILE2  .PUB  .TEST  DISC      :C         0      1 - 1      72    3   1
0    0
      FILE3  .PUB  .TEST  DISC      :C        32      1 - 1      72    1   2
111  112

```

Figure 6-5. illustrates the same example as above, except specifying the filenames in HFS syntax:

Figure 6-5. Long File Display (HFS Format)

FILENAME	VOLUME	RESTRICTIONS	SECTORS	CODE	MEDIA_NUM	RSIZE	BFCTR	#EXT	EOF	LIMIT
DISC	:C		1104		1 - 1	80	16	0	3500	3500
/TEST/PUB/FILE1										
DISC	:C		0		1 - 1	72	3	1	0	0
/TEST/PUB/FILE2										
DISC	:C		32		1 - 1	72	1	2	111	112
/TEST/PUB/FILE3										

If you use the `SHOW` parameter, but do not specify either `LONG` or `SHORT`, the system displays files in short form when the record size of the output device or file (`SYSLIST`) contains fewer than 132 characters in native mode and 114 characters in transport mode. The system displays files in long form when the record size is equal to or greater than 132 characters in native mode and equal to or greater than 114 characters in transport mode.

You can combine either `LONG` or `SHORT` with any of the other `SHOW` options, except `NAMESONLY`, but you cannot specify both `LONG` and `SHORT` at the same time within a `STORE` command.

Displaying File Dates

To display each file's creation date, last access date, and last modification date for the files stored, use the `SHOW=DATES` option. For example:

```
:FILE T;DEV=TAPE
:STORE @.PUB.DOC;*T;SHOW=DATES
```

Figure 6-6. illustrates a file listing created from the above example.

Figure 6-6. File Date Information (MPE Format)

FILENAME	GROUP	ACCOUNT	VOLUME	RESTRICTIONS	SECTORS	CODE	MEDIA
CREATED		ACCESSED	MODIFIED				
FILE1	.PUB	.TEST	DISC		:C	1104	1
1/12/1994		1/12/1994	1/12/1994				
FILE2	.PUB	.TEST	DISC		:C	0	1
1/12/1994		1/12/1994	1/12/1994				
FILE3	.PUB	.TEST	DISC		:C	32	1
1/12/1994		1/12/1994	1/12/1994				

When the output listing is in HFS format, an additional field is added to the dates display. The `STATE_CHANGE` field indicates the last date that a file attribute was changed. An example listing is:

Figure 6-7. illustrates a file listing created from the above example in HFS format.

Figure 6-7. File Date Information (HFS Format)

```

VOLUME RESTRICTIONS  SECTORS CODE MEDIA FILENAME
  CREATED      ACCESSED   MODIFIED   STATE_CHANGE

DISC           :C       1104     1         /TEST/PUB/FILE1
1/12/1994     1/12/1994  1/12/1994  1/12/1994

DISC           :C         0       1         /TEST/PUB/FILE2
1/12/1994     1/12/1994  1/12/1994  1/12/1994

DISC           :C        32       1         /TEST/PUB/FILE3
1/12/1994     1/12/1994  1/12/1994  1/12/1994

```

Displaying File Security Information

To display file security information for the files stored, use the `SHOW=SECURITY` option. The listing includes the file creator's user and account name, the security matrix, and a flag indicating the presence or absence of an ACD. The HFS format displays the same security information. For example:

```

:FILE T;DEV=TAPE
:STORE @.PUB.TEST;*T;SHOW=SECURITY

```

Figure 6-8. illustrates a file listing with security information from the above example.

Figure 6-8. File Security Information

```

FILENAME GROUP  ACCOUNT  VOLUME RESTRICTIONS  SECTORS CODE MEDIA
FILE1    .PUB  .TEST   DISC      :C      1104     1
LAURA  .STSUPPNM (R:ANY; A:ANY; W:ANY; L:ANY; X:ANY) *ACD EXISTS*

FILE2    .PUB  .TEST   DISC      :C         0       1
LAURA  .STSUPPNM (R:ANY; A:ANY; W:ANY; L:ANY; X:ANY) *ACD ABSENT*

FILE3    .PUB  .TEST   DISC      :C        32       1
LAURA  .STSUPPNM (R:ANY; A:ANY; W:ANY; L:ANY; X:ANY) *ACD ABSENT*

```

Displaying Names Only

The `SHOW=NAMESONLY` allows an abbreviated listing to be displayed. The only fields shown are the starting and ending media number, and the name of the file. This option is most useful when displaying HFS syntax files. Since they can be longer than MPE syntax names, it is sometimes useful to allow more room on the line to display the full pathname. `NAMESONLY` cannot be used with the `SHORT` or `LONG` options to `SHOW`. It has a slightly different format for MPE and HFS. The HFS style has the filename as the last field, while the MPE style has the filename as the first field.

Figure 6-9. is an example for HFS format.

Figure 6-9. Names Only Format

```
MEDIA_NUM  FILENAME
1 - 1      /_HFSACCT/_HFSCRIP/HFSMAP
1 - 1      /usr/lib/terminfo/e/ethernet
1 - 1      /usr/lib/terminfo/e/ex3000
1 - 1      /usr/lib/terminfo/e/exidy
1 - 1      /usr/lib/terminfo/e/exidy2500
```

Forcing the HFS Format

If SHOW=PATH is specified, the output STORE listing will be in HFS format, even if no HFS syntax files are being stored. You may want to use this so that all of your listings appear in the same format.

Displaying TurboSTORE/iX 7x24 True-Online Backup Information

When creating backups using the ONLINE=START or ONLINE=END options of the TurboSTORE/iX 7x24 True-Online Backup product, additional information is shown in the STORE listing. This additional information is a single character, immediately following the volume restrictions field. The values and meanings for this character are as follows:

- & After image file label data exists for this file, since the file label was modified before the sync point occurred. Only backups created with ONLINE=END will have files with after image file label data.
- # After image data exists for this file, since the file was modified before the sync point occurred. Only backups created with ONLINE=END will have after image file data.
- + This file was created during the backup, before the sync point occurred. Only backups created with ONLINE=END will have newly created files flagged this way.
- This file was deleted or renamed out of this backup before the sync point occurred. These files will not be restored by RESTORE. Backups created with ONLINE=START and ONLINE=END can have files flagged this way.

An example of these flags follows. Note that FOO5 was modified during the backup, FOO6's file label information was updated, FOO9 was removed from the backup, and FOO1 was added to the backup.

FILENAME	GROUP	ACCOUNT	VOLUME RESTRICTIONS	SECTORS	CODE	MEDIA
FOO2	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S	16	1
FOO3	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S	16	1
FOO4	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S	32	1
FOO5	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S#	16	1
FOO6	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S&	16	1
FOO7	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S	0	1
FOO8	.PUB	.MYACCNT	DISC	:C	32	1
FOO9	.PUB	.MYACCNT	MPEXL_SYSTEM_VOLUME_SET	:S-	0	1
FOO1	.PUB	.MYACCNT	DISC	:C+	0	1

Selecting an Error Recovery Method (ONERROR)

When STORE encounters an error, it either automatically recovers or terminates, depending upon the nature of the error and the error recovery method you choose.

Unrecoverable Errors

The following unrecoverable errors always cause STORE to terminate:

- A command syntax error.
- An error in the file system directory.
- An error opening a backup device or an indirect file. An indirect file is a text file containing STORE commands and options. Refer to "Using Indirect Files", in Chapter 5, "Storing Files".

Disk Read Errors

If STORE encounters a disk read error while storing a file, it skips the remainder of that file and sends an error message to the standard listing device describing the offset from the start of the file and the number of bytes that have been lost. STORE then continues to copy the remainder of the files. The STORE operation does not terminate. Since the files that contain disk errors are not stored, you will not be able to restore them.

Media Errors

When a media error occurs while writing to an output device, the behavior of STORE depends on the value of the ONERROR option, as explained below.

Tape Error Recovery The STORE command's ONERROR parameter lets you specify a backup error recovery procedure. Your options are ONERROR=QUIT and ONERROR=REDO. REDO is the default. QUIT instructs STORE to abort upon encountering a backup I/O error while REDO instructs STORE to continue after encountering an I/O error.

If you specify REDO and the STORE command encounters a backup I/O error on any device, STORE sends a message to your terminal (or the file or device to which you have assigned the file SYSLIST). This message indicates where the error occurred. For example, STORESET=(*T1 , *T2 , *T3) is requested and a backup error occurs on the device referenced by *T2 while it is creating reel 2. STORE rewinds the reel to the load point, marks the reel as bad, and asks you to mount another reel. After you mount the new reel, STORE automatically stores all the files from the point where the bad reel began.

The message is similar to the following:

- The user sees:

```
STORE ENCOUNTERED MEDIA WRITE ERROR ON LDEV #
```

where LDEV # is the device referenced by *T2. This means that the store error recovery occurs on the same device where the error occurred.

- The operator sees the following messages on the console:

```
STORE IS MARKING REEL BAD ON LDEV #
```

```
MOUNT MEDIA # OF SET # FOR STORE ON LDEV #
```

Optical Disk Errors If an optical disk write error is encountered, STORE displays the same error messages that it would display for tape backup devices.

You can determine which piece of media is has been marked bad by using the `STATUS` command of the `MOUTIL` program. The status display indicates which piece of media is bad by showing "THE MEDIA IS BAD" in the media name field. If a disk write error occurs while storing to a disk file, STORE aborts, irrespective of what `ONERROR` option was specified.

Bad Media

Once a piece of media has been marked bad, it should not be used for future STORE operations. If a media that has been marked bad is mounted for use in a STORE operation, STORE will issue the following message to the console:

```
MEDIA ON LDEV # IS BAD. CONTINUE ANYWAY?(Y/N)
```

If you reply YES, then STORE will write over this media. If you reply NO, the media will be unloaded, and you should mount another piece of media.

NOTE STORE will not check to see if a piece of media has been marked bad if it is the very first reel mounted. Therefore, it is important that once a reel has been marked bad by store, some note is made on the physical label of the media to indicate that it encountered a write error while storing to it.

Specifying Files by Date (DATE)

The `DATE` option specifies which files are stored based on two possibilities:

- STORE files that have been modified since a specified date
- STORE files that have not been accessed since a specified date

Storing Files Modified Since a Specified Date

You can use the `DATE` parameter to store only files that were modified or their states were changed since a specified date. (The state change date and time are updated any time the file is renamed or an ACD is changed. However, this information is not automatically shown. The state change date and time can be different from the last modified date displayed with the `LISTFILE filename,3` command. See the note below for information on displaying the state changes.)

The `DATE` parameter enables you to do, among other things, a partial system backup. For example, suppose a full system backup was performed on September 30, 1993. The following STORE command would copy all system and user files modified since that date:

```
:FILE T;DEV=TAPE  
:STORE @.@@;*T;DATE>=09/30/93
```

The date must be entered in either the above format (*mm/dd/yy*) or as *mm/dd/yyyy* (for example, 09/30/1993). The files stored include any modified on the same date as the date entered. The files stored also include those modified since the specified date on nonsystem volume sets that are online.

Use this form of the `DATE` parameter when you perform regular backups of a fileset. Store the entire set, for example, once a week and store only those files that were modified or their states were changed since the full backup the remaining six days.

NOTE STORE uses the state change date when storing files with the `DATE` parameter. To see the state change date, use the `FINFO(filename,42)` and `FINFO(filename,41)` functions.

The following command file displays the state change date and time for a specific file, such as `TSETUTIL` in this example:

```
parm file
setvar state_change 'STATE CHANGE = '+finfo('!file',42)+' , '+finfo('!file',41)
echo !state_change
deletevar state_change
```

Example:

```
(77)>chdate TSETUTIL
STATE CHANGE = TUE, OCT 31, 1995, 3:49 PM
(78)>
```

Storing Files Not Accessed Since a Specified Date

You can also use the `DATE` parameter to store only files that have not been accessed since a certain date. For example:

```
:FILE T;DEV=TAPE
:STORE @.OPERATOR.SYS;*T;DATE<=06/10/93
```

The year can be entered as two digits or four digits (93 or 1993). The above command stores any files in the `OPERATOR` group of the `SYS` account that have not been accessed since June 10, 1993. You may want to use this form of the `DATE` parameter with the `PURGE` parameter to remove out-of-date files from your system. Refer to the next section.

Removing Stored Files from the System: Archiving Files (PURGE)

If your system contains infrequently used or out-of-date files, you can store them onto backup and permanently remove them from the system disks at the same time. This process, called archiving files, increases available space and provides a backup copy of the files in case you need them in the future. To store disk files and then delete the disk files in a single operation, use the `PURGE` parameter of the `STORE` command. For example:

```
:FILE T;DEV=TAPE
:STORE @.OLD.ACCTG;*T;PURGE
```

The command above stores all files in the `OLD` group of the account `ACCTG` and deletes the files from disk. `PURGE` deletes the files only after it has successfully stored all of them. Thus, if `STORE` terminates in the middle of the process, it does not purge any files.

You can combine the `DATE` and `PURGE` parameters to periodically archive unused files. For example, the following `STORE` command stores to backup all files in the `EMPRECS` group of the `PERSONEL` account that have not been accessed since February 2, 1992 and then deletes the disk files:

STORE Options

```
:FILE T;DEV=TAPE
:STORE @.EMPRES.PERSONEL;*T;DATE<=02/02/92;PURGE
```

NOTE In order to use the `PURGE` parameter with a privileged file or files, you must have system manager (SM), system supervisor (OP), or privileged mode (PM) capability.

You may want to warn users before you archive files and inform them how to request to have archived files restored. If you need to restore purged files to the system, you can find them on the piece of media to which you stored them and on the media from the last full system backup you performed before you purged the files.

Displaying Progress Messages (PROGRESS)

Use the `PROGRESS` parameter to display `STORE` command progress messages at regular intervals. For example, to display progress messages every five minutes, use the following command:

```
:FILE T;DEV=TAPE
:STORE @.@@;*T;PROGRESS=5
```

If you use the `PROGRESS` parameter alone, without specifying an interval, `STORE` displays status messages every minute.

If `STORE` is being run from a MPE session, the progress messages will be sent to the standard list (`$STDLIST`). Otherwise, the messages will go to the system console.

System Directory Backup (DIRECTORY)

The system directory keeps track of the accounts, groups, users, and files in your system. Nonsystem volume set directories keep track of the contents of nonsystem volumes. Directory information can be stored when you back up system or nonsystem volumes. Storing directory information helps you rebuild your system from backups should it unexpectedly lose information. To store system directory information, use the `DIRECTORY` parameter. All HFS directories on the system are also stored.

The following example will store all files on the system, plus the system volume set accounting information.

```
:FILE T;DEV=TAPE
:STORE @.@@;*T;DIRECTORY
```

To store nonsystem volume set directory information, use `DIRECTORY` and `ONVS` in the same command. For example:

```
:FILE T;DEV=TAPE
:STORE @.@@;*T;ONVS=VOL_SET_A;DIRECTORY
```

In the above example, the `ONVS` parameter will limit `@@` to only files that reside on the specified volume sets. Only the directory information for groups and accounts on the specified volume sets will be stored. The system directory information will NOT be stored. See "Storing Files From Volume Sets" in this chapter for more information.

As illustrated above, the `DIRECTORY` parameter of the `STORE` command allows you to back

up the system and/or nonsystem volume set directories (account structure). You may then reload the account structure in the event of a system disaster by restoring the directory onto the system from the backup media. You have the added flexibility of backing up the directories of nonsystem volumes attached to the system as a part of or as a separate entity from the system. This allows a nonsystem volume (volume set) to become a separate reliable entity. System supervisor (OP) or system manager (SM) capability is required to use the `DIRECTORY` parameter.

NOTE When using the `DIRECTORY` parameter, the default fileset list is empty. This allows a backup to be created that contains only directory accounting information for the system volume set, or for the specified volume sets.

Here is an example of how to back up the system directory and all other files along with a nonsystem volume set directory and its files:

```
:FILE T;DEV=TAPE
:STORE @.@.*T;DIRECTORY;ONVS=MPEXL_SYSTEM_VOLUME_SET,NV1
```

where `NV1` is the name of a nonsystem volume set.

Storing Files in Compatible Formats (TRANSPORT)

The `TRANSPORT` option stores files on media that is compatible with:

- MPE V/E systems (`TRANSPORT`)
- Pre-POSIX systems (`TRANSPORT=MPEXL`)

Creating MPE V/E Compatible STORE Tapes

Use the `TRANSPORT` parameter of `STORE` to create `STORE` backups that are compatible with MPE V/E. For example, to store all of the files in the `PERSONEL` account onto an MPE V/E compatible backup, enter:

```
:FILE T;DEV=TAPE
:STORE @.@.PERSONEL;*T;TRANSPORT
```

The backup that is produced with the `TRANSPORT` parameter uses the same format as MPE V/E `STORE` backup format. Therefore, files that are not compatible with MPE V/E (due to file/record size limitations or new file types), are not stored if `TRANSPORT` is specified. The `TRANSPORT` option cannot be used with the following `STORE` and TurboSTORE options:

```
STORESET
INTER
DIRECTORY
ONVS
FCRANGE
SPLITVS
ONLINE
MAXTAPEBUF
COMPRESS
TREE
NOTREE
```

STORE Options

```

RENAME
NOTIFY
LOGVOLSET
PART[ IAL ]DB
FULLDB
STOREDIR[ ECTORY ]
NOSTOREDIRE[ ECTORY ]
STATISTICS

```

Character set specifications, as listed in "Using Wildcard Characters" in chapter 5, "Storing Files," cannot be used with the `TRANSPORT` parameter.

Creating Pre-POSIX Compatible Tapes

With the addition of HFS syntax files and other new POSIX-related features, mechanisms exist in `STORE` to allow backups to be created on a POSIX system and restored on a pre-POSIX system (4.0 or earlier). Pre-POSIX backups are created using the `TRANSPORT=MPEXL` option. For example:

```
:STORE @.MYGR.SOMEACCT;*T;SHOW;TRANSPORT=MPEXL
```

There are two general scenarios when files would need to be transferred to a pre-POSIX system:

- **Planned Transport**

The user stores off a set of files from one system, to be restored on another system right away.

- **Unplanned Transport**

The data stored off from a system (usually a system backup) is suddenly required to be restored on the same or a different system. This kind of transport is usually done to recover lost data and may require a POSIX system's backup to be restored on a pre-POSIX system.

With these two cases in mind, `STORE` has the following objectives with respect to data transport:

- The media created on a POSIX system should be restorable on a pre-POSIX as well as a post-POSIX release.
- The planned transport operation should be as simple and easy as possible.
- A default name conflict should not exist. This means that the naming of non-MPE syntax files on a transport to a pre-POSIX system should occur such that the user does not inadvertently overwrite previously transported files.
- For each of the transport situations, the security should be at least equal to or more restrictive than require on the target system.

The following sections list the features in `STORE` which address these objectives.

HFS Name Mapping Every POSIX `STORE` media is created such that any HFS syntax file can be restored to a pre-POSIX system with a unique MPE syntax name. When restoring a backup to a pre-POSIX system, the fileset of `@.@.HFSACCT` will contain all HFS syntax files on that media. These files are named with the MPE syntax of:

```
F#####._HFSGRP._HFSACCT
```

Where ##### represents numbered files starting at 0000001.

A mapping file is also on the media. This file specifies the mapping between the _HFSACCT name and the actual HFS name. The name of this file is HFSMAP._HFSGRP._HFSACCT.

Sample entries in the file would look similar to this:

```
F0000001._HFSGRP._HFSACCT <- /SYS/PUB/dir1/file_a
F0000002._HFSGRP._HFSACCT <- /SYS/PUB/dir1/file_b
F0000003._HFSGRP._HFSACCT <- /SYS/PUB/dir1/file_c
```

Since _HFSGRP and _HFSACCT are illegal MPE filenames, special steps must be taken when attempting to restore these files on a pre-POSIX system. The LOCAL option can be used to restore the HFS files into your logon group. For example, on a pre-POSIX system, issue:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@.@HFSACCT;SHOW;LOCAL
```

Alternately, you can specify a specific group and account to place the HFS files, using the CREATE, GROUP, and ACCOUNT options:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@.@HFSACCT;SHOW;CREATE;GROUP=NEWGRP;ACCT=MYACCT
```

This will restore all of the HFS syntax files to the group NEWGRP.MYACCT. The HFSMAP file will allow you to identify the real names of the F##### files.

If you want to only restore MPE syntax files from the backup, that can be accomplished by the following command:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@.@ - @.@.@HFSACCT;SHOW;CREATE
```

NOTE There are certain types of files that exist on a POSIX system that have no equivalent on a pre-POSIX system, such as symbolic links and device links. It is not recommended that you attempt to restore these files on a pre-POSIX system. They will contain no useful information for that system, and may be difficult to use or purge.

Preserving File Security An argument of the TRANSPORT option allows POSIX style ACD's to be translated to pre-POSIX systems, while maintaining the correct security. Using TRANSPORT=MPEXL will cause STORE to write out the pre-POSIX format ACD's. If a POSIX media is created without TRANSPORT=MPEXL, and then restored to a pre-POSIX system, the ACD's will not be understood and the system will enforce a "creator only" security. This provides a more restrictive security. Also, all POSIX-related fields of the file labels being stored are cleared on the backup media.

Specifying the Number of Files to Store (FILES)

The FILES parameter works with the TRANSPORT parameter. By default, MPE/iX stores a maximum of 4000 files at a time. If you are storing more than 4000 files using the TRANSPORT parameter, use the FILES= parameter to specify the maximum number of files to be stored. If you do not use the TRANSPORT parameter, the FILES= parameter is ignored. For example, to set the maximum number of files stored to 6000 for a TRANSPORT backup,

you might enter the following:

```
:FILE T;DEV=TAPE  
:STORE @.OPERATOR.SYS;*T;TRANSPORT;FILES=6000
```

Use the MPE/iX `RESTORE` command to restore files from MPE V/E format backup tapes to an MPE/iX system. MPE/iX `RESTORE` automatically detects that the tapes are in MPE V/E format, and invokes `CM(RE)STORE` (Compatibility Mode Restore) to restore the files.

Managing File Security (COPYACD and NOACD)

An additional way to manage the security of files being stored is to use the `COPYACD` and `NOACD` options. The default is `COPYACD`. This means that in addition to the MPE security matrix, a file's ACD (if one exists) is written to backup. This allows you to store the extra security information associated with an ACD.

If for some reason you want to store files without their ACD, the `NOACD` option can be specified. This would remove the extra security associated with ACD's, which may be desired if you are moving files to a system where the current ACD's would be inapplicable or undesired. When using these options, keep in mind that `RESTORE` also has `COPYACD` and `NOACD` options. So, even if you create a `STORE` media with ACD's, you have the choice at `RESTORE` time to restore or not restore the ACD information.

Storing Files with Certain File Codes (FCRANGE)

Use the `FCRANGE` parameter to store files with certain file codes. MPE XL file codes distinguish different types of files. You can select up to eight file code ranges to store. For example, files with codes 1100, 1101, and 1102 are HPWORD files. Files with codes 1152 and 1153 are SLATE files. To store all HPWORD and SLATE files in the `LETTERS` group of the `PERSONEL` account enter:

```
:FILE T;DEV=TAPE  
:STORE @.LETTERS.PERSONEL;*T;FCRANGE=1100/1102,1152/1153
```

Improving Storing Efficiency by Increasing Record Size (MAXTAPEBUF)

Files are stored with greater efficiency by using the maximum backup record (buffer) size allocated to a given backup device. Performance increases because data transmission is more efficient and because the amount of media used for interrecord gaps decreases.

The maximum backup record size of MPE/iX supported backup devices differs from device to device and is shown in Table 6-1.

Table 6-1. Maximum Tape Record Sizes

Tape Drive	Record Size
HP 7974	16 KB
HP 7978A	16 KB
HP 7978B	32 KB
HP 7979	32 KB
HP 7980(XC,S,SX)	32 KB
DDS (all)	32 KB
3480	32KB

MAXTAPEBUF should be used when storing to magnetic tape, DDS, or 3480 cartridges. Currently, it is not necessary to specify MAXTAPEBUF when storing to MO devices or disk files, since the maximum record size of 32K is already being used.

Normally when a STORE command is issued, a default record size of 16 kilobytes is used. However, by using the MAXTAPEBUF parameter of the STORE command, the maximum record size of the specified device is automatically used. If multiple backup devices are specified, using the STORESET parameter, in conjunction with MAXTAPEBUF, STORE uses the maximum record size, which is the lowest common denominator of the different devices. This is necessary to ensure that the record sizes are the same on all volumes of the backup set.

NOTE Do not use the MAXTAPEBUF parameter with a backup device of 32 KB record size if you anticipate that you will later restore the backup created to a backup device whose maximum record size is 16 KB. Restoring a backup to a larger maximum record size device, on the other hand, does not cause any problem.

To increase record size when storing files, use the MAXTAPEBUF parameter in your STORE command. For example:

```
FILE T;DEV=TAPE
STORE @.OPERATOR.SYS;*T;MAXTAPEBUF
```

Notification of File Availability (NOTIFY)

During the STORE process, the fileset that you are storing is usually unavailable for modification. Files can be modified again after STORE is completed. If you are doing an on-line backup using:

- TurboSTORE/iX 7x24 True-Online Backup *online backup* (using the ONLINE option only), the files being stored are unavailable for a short time at the beginning of the store. However, once they have been attached to STORE's logging mechanism, they are available for modification.
- TurboSTORE/iX 7x24 True-Online Backup with *true-online backup*, (using the

STORE Options

ONLINE option with START or END), the files being stored remain available and can be modified while the backup is occurring. Refer to Chapter 7, "TurboSTORE/iX 7x24 True-Online Backup," for more information.

The NOTIFY option can be used to help automate the process of performing some action at the time where the files being stored are available for modification, either at the end of the STORE, at the end of the attach phase, or at the end of the sync point.

NOTIFY allows STORE to stream a user-specified job file (using the MPE/iX STREAM command) either at the end of the attach phase (while using ONLINE), at the end of a successful STORE operation (normal backup), or at the end of the sync point (ONLINE=START or ONLINE=END). The job file is specified by the user by setting up a formal file designator named NOTIFY, prior to running the store job. This job can contain various actions that notify users that the files are free, start other jobs that modify the files, or allow users back on to the system, for example.

If the specified job does not contain passwords, STORE will prompt the user for the passwords if STORE is being run from a session. If STORE is being run from a job and passwords are not present, the job stream will fail.

When the job fails to stream, due to lack of passwords, or other reasons, all error messages will be sent to the standard list (\$STDLIST). All other output as the result of streaming the job will also be sent to \$STDLIST.

The following is an example of using the NOTIFY command in conjunction with a user specified job file.

First, set up the file equation, using the FILE command:

```
:FILE NOTIFY=MYJOB.PUB.SYS
```

Where the job stream could contain, for example:

```
!JOB NOTIFY,MANAGER.SYS
!TELL @ Your files are now available!
!EOJ
```

Then, run STORE with the NOTIFY option:

```
:FILE T;DEV=TAPE
:STORE @.PUB.SYS;*T;SHOW;ONLINE;NOTIFY
```

Since this is an online store, the job stream will be executed once the files have been attached, and a message would then be sent to all users telling them that their files are now free to be accessed.

Additional Backup Information (STATISTICS)

In addition to the counts of files stored and not stored that are normally reported at the end of a backup, the STATISTICS option provides extra data about the backup. This extra information includes:

- Amount of data written to each piece of media for each parallel set
- Amount of time required to write each piece of media
- Throughput for each piece of media

- Number of retries incurred for each piece of media

If software compression is enabled with the `COMPRESS` option, additional data on the amount of compressed data written, and compression ratios and percentages is displayed.

If an online or 7x24 true-online backup is performed, then information on the amount of log data written is displayed.

All data totals are shown in bytes. To convert these totals to sectors, divide by 256.

When displaying the amount of backup data written, up to five columns of data can be displayed, depending on what options were specified for the backup. These columns display the following backup data information:

File Data	This is the amount of the backup that comes from the actual data in files. Most of the data in your backup will be accounted for in this column. If you are using software compression (<code>COMPRESS=LOW</code> or <code>COMPRESS=HIGH</code>), then this is the total number of <i>uncompressed</i> bytes of file data written. The actual, compressed amount of data written is displayed in the "Compressed" column.
Misc Data	This is data that STORE needs to write to the backup to manage what files are on the backup, such as the STORE backup label and a media directory. The amount of miscellaneous data on a backup is proportional to the number, but not size, of files being stored.
Compressed	This column is only displayed if software compression was used to create the backup (with <code>COMPRESS=LOW</code> or <code>COMPRESS=HIGH</code>). It displays the total number of compressed bytes of file data written to the backup. The ratio between this column and the file data column is used to display the Compression ratios and percentages later in the statistics listing.
Log Data	If an online backup (<code>ONLINE</code>) or 7x24 true-onlinebackup (<code>ONLINE=START</code> or <code>ONLINE=END</code>) backup is performed, this column displays how much log data was written to the backup. This includes both before image and after image log data. (For a definition of before and after image log data, refer to Chapter 7). This total depends on how much activity occurred on the files being backed up during the backup. This is a good total to check if your online backup is taking longer or consuming more media than you think it should. Such a slowdown could be attributed to too much activity on files being stored, which slows down both the users and the online backup.
Total on Reel	This column is a sum of all the other columns displayed for this particular piece of media. It represents the entire amount of data stored for that media. Note that if software compression was used, then the Compressed column, and NOT the File Data column, is used to compute the media total. If no software compression is used, then the File Data column is used.

These data columns are displayed for each piece of media written to during the backup. If a parallel backup is performed (using the `STORESET` option), then the media will be grouped by parallel set. After all parallel sets are displayed, totals for each parallel set, and then for the entire backup are displayed. Look for the "Backup Total" line to determine

the total amount of all data written for the backup.

The Throughput section follows the data written columns. The throughput for a particular piece of media is computed by dividing the total amount of data written for a piece of media by the time it took to completely write that data. The time required to switch media is NOT accounted for in this computation. Only time spent actually writing the media is used. Also, note that if software compression is being used, then the amount of compressed data written is used to compute the throughput. Your *virtual* throughput may actually be higher, if software compression is used.

The number of hardware retries encountered when writing the data is also displayed here. This only applied to magnetic tape, DDS, and 3480 media. It has no meaning when writing to disk files or magneto-optical devices, and will always be zero. Monitoring the number of retries for a piece of media is important in determining if that media is starting to go bad. An increased number of retries could indicate that the media should no longer be used, or that your devices should be cleaned and/or serviced.

After the Throughput section, the total elapsed time for the backup is displayed. This time does include media switches.

If software compression is used for the backup (with COMPRESS=LOW or COMPRESS=HIGH), then a section on compression ratios will follow the Throughput section. For each media, or each media in each parallel set, the amount of uncompressed file data written is displayed with the amount of data actually written when that file data is compressed. The ratio between these two numbers and a compression percentage is also displayed. The higher the ratio is, the better the data compressed. Similarly, the higher the percentage, the better the data compressed. A negative percentage is possible, and indicates that the data you were storing did not compress at all. If you see this in your statistics display, or see very low compression ratios, it indicates that you are storing data that does not compress well, and that it might be better not to use software compression, or try a different software compression (LOW instead of HIGH, for example).

NOTE These compression ratios are highly dependent on the type of data being stored, and may vary from the Hewlett-Packard supplied compression ratios. In general, data files will compress better than application and system files. You should carefully pick the type of data compression used to match the type of data being stored.

Following are examples of output from the STATISTICS options.

Example 1

The following example is of a normal backup to a single device without software compression or online backup.

```

BACKUP STATISTICS

DATA WRITTEN

      FILE DATA      MISC DATA      TOTAL ON REEL
MEDIA 1 | 189826798 | 1389164 | 191215962
-----+-----+-----+-----
BACKUP
TOTAL      189826798      1389164      191215962

THROUGHPUT
      MEDIA START TIME      MEDIA END TIME      THROUGHPUT (K/S)      RETRIES
MEDIA 1      11:45 AM      12:02 PM      177 K/S      0

TOTAL ELAPSED TIME: 0 hours, 17 minutes, 31 seconds

```

Example 2

The following example is of a backup to two pieces of media.

```

BACKUP STATISTICS

DATA WRITTEN

      FILE DATA      MISC DATA      TOTAL ON REEL
MEDIA 1 | 13538720 | 166480 | 13705200
MEDIA 2 | 176288078 | 747888 | 177035966
-----+-----+-----+-----
BACKUP
TOTAL      189826798      914368      190741166

THROUGHPUT
      MEDIA START TIME      MEDIA END TIME      THROUGHPUT (K/S)      RETRIES
MEDIA 1      11:41 AM      11:41 AM      669 K/S      0
MEDIA 2      11:41 AM      11:45 AM      835 K/S      0

TOTAL ELAPSED TIME: 0 hours, 3 minutes, 49 seconds

```

Example 3

The following example is of a backup to parallel devices with software compression and online backup.

STORE and TurboSTORE/iX Command Options
STORE Options

BACKUP STATISTICS

DATA WRITTEN

SET 1	FILE DATA	MISC DATA	COMPRESSED	LOG DATA	TOTAL ON REEL
MEDIA 1	63344322	418976	36349232	0	36768208

SET 1					
TOTAL	63344322	418976	36349232	0	36768208
SET 2	FILE DATA	MISC DATA	COMPRESSED	LOG DATA	TOTAL ON REEL
MEDIA 1	66652980	307716	22195808	6144	22509668

SET 2					
TOTAL	66652980	307716	22195808	6144	22509668
SET 3	FILE DATA	MISC DATA	COMPRESSED	LOG DATA	TOTAL ON REEL
MEDIA 1	59869176	164384	8199956	0	8364340

SET 3					
TOTAL	59869176	164384	8199956	0	8364340

SET TOTALS

SET 1	63344322	418976	36349232	0	36768208
SET 2	66652980	307716	22195808	6144	22509668
SET 3	59869176	164384	8199956	0	8364340

BACKUP

TOTAL	189866478	891076	66744996	6144	67642216
-------	-----------	--------	----------	------	----------

THROUGHPUT

	MEDIA	START TIME	END TIME	THROUGHPUT (K/S)	RETRIES
SET 1					
MEDIA 1	12:15 PM	12:19 PM	147 K/S	0	
SET 2					
MEDIA 1	12:15 PM	12:18 PM	100 K/S	0	
SET 3					
MEDIA 1	12:15 PM	12:17 PM	50 K/S	0	

TOTAL ELAPSED TIME: 0 hours, 4 minutes, 3 seconds

COMPRESSION

	FILE DATA	COMPRESSED	RATIO	PERCENTAGE
SET 1				
MEDIA 1	63344322	36349232	1.7:1	43%
SET 2				
MEDIA 1	66652980	22195808	3.0:1	67%
SET 3				
MEDIA 1	59869176	8199956	7.3:1	86%
BACKUP TOTAL	189866478	66744996	2.8:1	65%

Storing Files from Volume Sets (ONVS)

To store the files on a particular volume set, use the ONVS parameter of the STORE command. ONVS lets you store files from up to 20 volume sets at once. For example, to store all files from the nonsystem volume sets named VOL_SET_A and VOL_SET_B, enter:

```
:FILE T;DEV=TAPE
:STORE @.@@;*T;ONVS=VOL_SET_A,VOL_SET_B
```

To store files from system and nonsystem volumes, include the system volume set name within the ONVS parameter. For example:

```
:FILE T;DEV=TAPE
:STORE
@.@@;*T;ONVS=MPEXL_SYSTEM_VOLUME_SET,&
VOL_SET_A,VOL_SET_B
```

NOTE Although the ONVS and SPLITVS options can be used together, a volume set that is specified for ONVS cannot be specified for SPLITVS and a volume set that is specified for SPLITVS cannot be specified for ONVS.

Storing Files from Split Volumes (SPLITVS)

Use the SPLITVS option to store files that are located on the specified split volume sets. These volumes are also known as mirrored disks.

This provides a form of concurrent, or online backup, because you can:

- Split a volume set.
- Allow users to access their files
- Perform a backup on the mirrored volume

The SPLITVS option requires:

- The Mirrored Disk/iX product be installed on the system.
- The specified volume sets must be configured as mirrored volumes.

The basic procedure to use split volumes and STORE together is:

1. Request users to close the files on the volume sets that are to be backed up for a short duration.
2. VSCLOSE the mirrored volume set with the SPLIT option.
3. VSOOPEN the volume set. This will open the user and backup sets.
4. Allow users to begin accessing the files again.
5. Issue the store command, using the SPLITVS option. For example:

```
:STORE @.@@;*T;SHOW;SPLITVS=my_split_set
```

6. After the backup has completed, join the split set back together using the JOINMIRRSET command in VOLUTIL.

For more information on using mirrored disks, see the *Mirrored Disk/iX User's Guide*.

NOTE Although the `ONVS` and `SPLITVS` options can be used together, a volume set that is specified for `ONVS` cannot be specified for `SPLITVS` and a volume set that is specified for `SPLITVS` cannot be specified for `ONVS`.

Renaming Files (RENAME)

When using the `RENAME` option, an original file refers to the original disk file, and a target file indicates the final filename on the backup media.

Renaming a file requires the capabilities necessary to manipulate files across groups and accounts. Table 6-1. assigns capabilities to the different levels of file renaming.

Table 6-2. Required Capabilities for RENAME

Capabilities Required	File	Group	Acct	Creator
File Creator	Yes	Yes	No	No
Account Manager (AM)	Yes	Yes	No	Yes
System Manager (SM)	Yes	Yes	Yes	Yes
System Supervisor (OP)	Yes	Yes	Yes	Yes

The `RENAME` option has the following syntax:

```
STORE file set [=targetname [:creatorname [.creatoraccount] ]] [,file  
set[=targetname [:creatorname [.creatoraccount] ] ]] ;*TAPE;RENAME;SHOW
```

The fileset consists of the inclusion fileset and all exclusion filesets. All legal filesets are allowed. The targetname can be any legal MPE or HFS filename with the following restrictions on wildcards:

1. Each component of the file or pathname must either be equal to `@` or have no wildcards.
2. An HFS pathname which ends in a `/` can have no wildcards. The trailing slash is used to designate a directory to place all files.

Each component of the creator name must either be equal to `@` or have no wildcards. The account part of the creator name may be omitted.

The `@` is used to specify that the component in the original filename should be used for that component of the target filename. When specifying HFS pathnames, components will be replaced from the root on down. However, if the last component of the target pathname is an `@` the last component of the original pathname will be used. If there is not a corresponding component for an `@` in the original pathname the file will not be stored.

The following tables illustrate sample translations for the target filenames:

STORE a.b.c=...;TAPE;RENAME;SHOW

Original Filename	Target Specification	Target Filename
A.B.C	x.y.z	X.Y.Z
A.B.C	@.y.z	A.Y.Z
A.B.C	x.@.z	X.B.Z
A.B.C	x.y.@	X.Y.C

STORE /A/B/C=...;TAPE;RENAME;SHOW

Original Filename	Target Specification	Target Filename
/A/B/C	@.y.z	C.Y.Z
/A/B/C	x.@.z	X.B.Z
/A/B/C	x.y.@	X.Y.A

STORE /a/b/c=...;TAPE;RENAME;SHOW

Original Filename	Target Specification	Target Filename
/a/b/c	@.y.z	/z/y/c
/a/b/c	x.@.z	/z/b/x
/a/b/c	x.y.@	/a/y/x

STORE /a=@/x/@/y/@;*TAPE;RENAME;TREE;SHOW

STORE Options

Original Filename	Target Specification	Target Filename
/a/	/@/x/@/y/@	NOT STORED: *
/a/f1	/@/x/@/y/@	NOT STORED: *
/a/b/	/@/x/@/y/@	NOT STORED: *
/a/b/f2	/@/x/@/y/@	NOT STORED: *
/a/b/c/	/@/x/@/y/@	NOT STORED: *
/a/b/c/f3	/@/x/@/y/@	/a/x/c/y/f3
/a/b/c/d/	/@/x/@/y/@	/a/x/c/y/d/
/a/b/c/d/f4	/@/x/@/y/@	/a/x/c/y/f4
/a/b/c/d/e/	/@/x/@/y/@	/a/x/c/y/e/
/a/b/c/d/e/f5	/@/x/@/y/@	/a/x/c/y/f5
/a/b/c/d/e/f/	/@/x/@/y/@	/a/x/c/y/f/
/a/b/c/d/e/f/f6	/@/x/@/y/@	/a/x/c/y/f6
/a/b/c/d/e/f/g/	/@/x/@/y/@	/a/x/c/y/g/
/a/b/c/d/e/f/g/f7	/@/x/@/y/@	/a/x/c/y/f7

*Source name does not have component for substitution

The following table illustrates sample translations for a target file directory when the fileset specification was not a directory:

```
STORE /a/b=/w/x/y/z/ ; *TAPE ; RENAME ; TREE ; SHOW
```

Original Filename	Target Specification	Target Filename
/a/b/	/w/x/y/z/	/w/x/y/z/b/
/a/b/f2	/w/x/y/z/	/w/x/y/z/b/f2
/a/b/c/	/w/x/y/z/	/w/x/y/z/c/
/a/b/c/f3	/w/x/y/z/	/w/x/y/z/c/f3
/a/b/c/d/	/w/x/y/z/	/w/x/y/z/d/
/a/b/c/d/f4	/w/x/y/z/	/w/x/y/z/d/f4
/a/b/c/d/e/	/w/x/y/z/	/w/x/y/z/e/
/a/b/c/d/e/f5	/w/x/y/z/	/w/x/y/z/e/f5
/a/b/c/d/e/f/	/w/x/y/z/	/w/x/y/z/f/
/a/b/c/d/e/f/f6	/w/x/y/z/	/w/x/y/z/f/f6
/a/b/c/d/e/f/g/	/w/x/y/z/	/w/x/y/z/g/
/a/b/c/d/e/f/g/f7	/w/x/y/z/	/w/x/y/z/g/f7

The following table illustrates sample translations for a target file directory when the fileset specification was a directory:

```
STORE /a/b/= /w/x/y/z/ ; *TAPE ; RENAME ; SHOW
```

Original Filename	Target Specification	Target Filename
/a/b/	/w/x/y/z/	/w/x/y/z/
/a/b/f2	/w/x/y/z/	/w/x/y/z/f2
/a/b/c/	/w/x/y/z/	/w/x/y/z/c/
/a/b/c/f3	/w/x/y/z/	/w/x/y/z/c/f3
/a/b/c/d/	/w/x/y/z/	/w/x/y/z/c/d/
/a/b/c/d/f4	/w/x/y/z/	/w/x/y/z/c/d/f4
/a/b/c/d/e/	/w/x/y/z/	/w/x/y/z/c/d/e/
/a/b/c/d/e/f5	/w/x/y/z/	/w/x/y/z/c/d/e/f5
/a/b/c/d/e/f/	/w/x/y/z/	/w/x/y/z/c/d/e/f/
/a/b/c/d/e/f/f6	/w/x/y/z/	/w/x/y/z/c/d/e/f/f6
/a/b/c/d/e/f/g/	/w/x/y/z/	/w/x/y/z/c/d/e/f/g/
/a/b/c/d/e/f/g/f7	/w/x/y/z/	/w/x/y/z/c/d/e/f/g/f7

In a similar manner, the wildcard @ is used to specify the use of the original creator name component in the new creator name. If no creator is specified, the original creator is kept. This is equivalent to specifying @.@ in the creator field. If the creator name is specified without an account part, the account part of the creator name will be made equal to the account of the new file.

There is an exception to this rule when the new filename does not have a legal account name in its pathname. A directory off of the root node that has a legal account name will be considered the account for the creator. This is true whether this directory is an account or not.

Table 6-1. below illustrates the translations for the creator. All files below in the "sg" group of the "sa" account were created by "CREATOR.CRACCT."

Table 6-3. Creator File Translations

Original Filename	Target Specification	Target Creator
s1.sg.sa	d1.dg.da	CREATOR.DA
s2.sg.sa	d2.dg.da:@.@	CREATOR.CRACCT
s3.sg.sa	d3.dg.da:user.acct	USER.ACCT
s4.sg.sa	d4.dg.da:user	USER.DA
s5.sg.sa	d5.dg.da:user.@"	USER.CRACCT
s6.sg.sa	d6.dg.da:@.acct	CREATOR.ACCT
s7.sg.sa	d7.dg.da:@	CREATOR.DA
s8.sg.sa	/d8	CREATOR.CRACCT
s9.sg.sa	/d9:@.@	CREATOR.CRACCT
s10.sg.sa	/d10:user.acct	USER.ACCT
s11.sg.sa	/d11:user	USER.CRACCT
s12.sg.sa	/d12:user.@"	USER.CRACCT
s13.sg.sa	/d13:@.acct	CREATOR.ACCT
s14.sg.sa	/d14:@	CREATOR.CRACCT
s15.sg.sa	/ACCOUNT/GROUP/D15	CREATOR.CRACCT
s16.sg.sa	/ACCOUNT/GROUP/D16:@.@	CREATOR.CRACCT
s17.sg.sa	/ACCOUNT/GROUP/D17:user.acct	USER.ACCT
s18.sg.sa	/ACCOUNT/GROUP/D18:user	USER.CRACCT
s19.sg.sa	/ACCOUNT/GROUP/D19:user.@"	USER.CRACCT
s20.sg.sa	/ACCOUNT/GROUP/D20:@.acct	CREATOR.ACCT
s21.sg.sa	/ACCOUNT/GROUP/D2:@1	CREATOR.CRACCT

If the SHOW option is specified, only the target names are listed. The security option of the SHOW option displays the target creator name if one is specified.

The ACD associated with a file is not changed when the file is renamed. As with other STORE operations, the ACD is copied unless the NOACD option is specified. If you need to change the ACD, use the ALTSEC command.

You can specify the same target file information for more than one org file. However, when two or more files have the same name on a backup created using the `RENAME` option, `RESTORE` will write the second file over the first file.

The files excluded may only be specified on the original side of the fileset. As with normal exclusion, there is no limit to the number of excluded filesets. The example below shows an exclusion file specification:

```
:STORE @.PUB.SYS -a -b -c -d=@.backup.sys:creator;*t;rename
```

Lockwords are not allowed on target files, and if specified cause `STORE` to abort. The options that are not supported with the `RENAME` option and will cause `STORE` to abort are:

- PURGE
- TRANSPORT
- SPLITVS

Specifying HFS Files (TREE and NOTREE)

By default, `STORE` uses the last character of a HFS filename to determine if the specified file or fileset should be scanned recursively to include all files below the specified wildcards or directories. If a filename ends in a `/`, then it will be scanned recursively to include all files below it in the hierarchical directory. Otherwise, the file or fileset is not scanned recursively, resulting in a horizontal cut at its level in the hierarchical directory. Using `TREE` and `NOTREE` can override that default behavior.

If `TREE` is specified, ALL files and filesets are scanned recursively, regardless of their ending character. If `NOTREE` is specified, then NO files or filesets are recursively scanned, resulting in a horizontal cut in the directory structure.

Storing Database Files Using PARTIALDB or FULLDB

By using the new options `PARTIALDB` or `FULLDB`, you can back up either a partial or an entire TurboIMAGE database and/or ALLBASE/SQL DBEnvironment. The type of backup you are performing determines the type of option you would use:

- `PARTIALDB` can be used with a 7x24 true-online backup (using `ONLINE=START` or `ONLINE=END`)
- `FULLDB` can be used with a non-7x24 true-online backup (all other backups)

Storing Database Files Using 7x24 True-Online Backup

When performing a 7x24 true-online backup using the `ONLINE=START` or `ONLINE=END` options, by default, `STORE` backs up the entire database.

- To back up a TurboIMAGE database, specify the name of the root file only in the fileset list.
- To back up an ALLBASE/SQL DBEnvironment, specify the name of the DBE file.

Specification of the TurboIMAGE root file or the ALLBASE/SQL DBE file causes `STORE` to store the entire database. If a dataset file is specified without the root file, then `STORE` prints the following warning:

STORE Options

```
PARTS22.GROUP.ACCT NOT STORED: FILE IS PART OF AN IMAGE
DATABASE AND ROOT IS NOT SPECIFIED
```

If any of the dataset files are specified in addition to the root file, no warning will be displayed but the individual dataset files will be counted as redundantly specified files.

If you want to back up a partial database, you can use the `PARTIALDB` option. STORE will backup only those dataset files specified in the fileset list.

CAUTION When using the `PARTIALDB` option, if the root file is missing from the fileset list, the quiesce may not succeed (refer to Chapter 7, "TurboSTORE/iX 7x24 True-Online Backup," for more information on quiescing). STORE will still back up the dataset files listed in the fileset list.

Since quiescing a partial database cannot be guaranteed, we recommend that `PARTIALDB` **not be used** with a 7x24 true-online backup.

NOTE Quiescing an ALLBASE/SQL database and storing the ALLBASE/SQL database by TurboSTORE/iX 7x24 True-Online Backup when the DBEnvironment is in normal use, has the same effect as using the `STOREONLINE` command of SQLUtil. For a 7x24 true-online backup of ALLBASE/SQL to be of use for rollforward recovery, use the same procedures you would use when doing the `STOREONLINE` commands.

Storing Database Files Using Non-7x24 True-Online Backup

When performing a non-7x24 true-online backup (where you do not specify the 7x24 true-online backup options, `ONLINE=START` or `ONLINE=END`), by default, STORE allows a partial database backup of only those database files included in the fileset list. For example:

```
STORE parts01,parts02;... Stores only database files parts01 and parts02.
STORE parts@;... Stores entire database parts as well as any files starting
with parts including positive and negative filecodes.
STORE parts;... Stores only the TurboIMAGE root file parts or the
ALLBASE/SQL DBE file partsdbe.
```

If you want to store an entire database, use the new option, `FULLDB`.

Creating a Disk Directory File (STOREDIRECTORY and NOSTOREDIRECTORY)

To increase the ease and speed of file recovery, all STORE and TurboSTORE products have the ability to put a copy of a backup's STORE label and STORE directory into a disk file. This directory information allows RESTORE to more quickly determine where files are located on the backup, and to prompt users to mount the necessary media. If you want to enable this feature, specify the `STOREDIR[ECTORY]` option.

NOTE This information in this file is also used to determine if a file being restored has after image log data. A file will have after image log data if it was modified before the sync point during a 7x24 true-online backup. RESTORE must have this information before restoring files, so when creating a 7x24 true-online backup with the sync point at the end of the backup, the STOREDIR option is automatically enabled. See Chapters 7 and 11 for more information on creating and restoring 7x24 true-online backups.

The filename used by STORE uniquely identifies the file and associates it with a specific backup. The file is placed in the store_dirs HFS directory in the HPSTORE.SYS group.

WARNING The HPSTORE group and store_dirs directory do NOT exist on the system by default. You must create them, and set up the appropriate security before using the STOREDIR option. STORE will NOT print a warning if it cannot create the file because the path does not exist. See the following paragraphs for information on creating and maintaining the HPSTORE group and store_dirs directory.

The following naming convention is used:

```
store_yyyymmdd_hhmmssstt_pin##_day
```

For example:

```
/SYS/HPSTORE/store_dirs/store_19940622_13300314_pin30_wed
```

This name contains the date and time of the backup and the pin of the STORE process. This naming convention allows a LISTFILE, 2 to show a chronological list of STORE directory files.

You should not rename or purge STORE directory files that you plan on using, since RESTORE looks for a specifically named file. However, if the file does not exist, all files can still be restored from the backup. The size of the directory file is directly proportional to the number of files stored in the backup. You can control how much disk space is used for these directory files by putting a disk file limit on HPSTORE.SYS.

The following example limits the space used by the HPSTORE group to 4000 sectors:

```
:ALTGROUP HPSTORE.SYS;FILES=4000
```

You can have STORE associate a shorter, more meaningful name to the disk directory files by specifying an optional filename argument to STOREDIR. You can specify an MPE or HFS filename. STORE then creates a symbolic link with the specified name, which points to the original disk file.

For example, to create a full backup on Friday and give it the name FRIFULL, issue the command:

```
:STORE /*tape;STOREDIRECTORY=FRIFULL.HPSTORE.SYS
```

To create a partial backup of all files that were modified since May 4, 1994, and create a disk directory filename, WedPartial, in the current working directory, issue the command:

```
:STORE /*tape;STOREDIR=./WedPartial;DATE>=5/4/94
```

If you want to create a 7x24 true-online backup with the sync point at the end of the backup, and **do not** want a disk directory file to be created, specify the `NOSTOREDIRECTORY` or `NOSTOREDİR` option on the `STORE` command line.

Storing Files in Interleave Format (INTER)

With `STORE`, `TurboSTORE/iX II`, and `TurboSTORE/iX 7x24 True-Online Backup`, you have the capability of simultaneously reading information from different files on multiple disk drives and storing the accessed information on the backup device in a single operation.

This process called interleaving allows you to match slow input devices (such as disk drives), with fast output devices (such as streaming tape devices). This effectively increases the input rate.

By varying the number of concurrent input operations, the effective input rate can be matched with the output rate. This allows the backup to execute closer to the speed of the output device.

When you `STORE` files in interleave format, `STORE`:

1. Partitions the set of files to your backup device or devices.
2. Opens the files. It may (depending upon file characteristics), access several files at the same time.
3. Interleaves files on the same backup to make sure that your backup device or devices operate at the highest speed possible.

Interleave format significantly reduces the time necessary to store a large set of files, such as when performing a full system backup. It merges the stream of data from multiple files on separate disks into a single stream of data sent to the backup device or devices. If necessary, `STORE` processes files out of sequence to keep backup devices operating efficiently.

Interleave format is most efficient when your fileset is evenly spread across two or more system disks. Use sequential devices or parallel devices alone when the set of files to store is smaller and less evenly distributed, such as when doing a partial backup, or if the files are only coming from one disk.

To `STORE` files in interleave format, use the `INTER` parameter. For example:

```
:FILE T;DEV=TAPE  
:STORE @.OPERATOR.SYS;*T;INTER
```

When these files are restored, `RESTORE` will automatically recognize the interleave format and restore the files properly.

Using the `STORESET` and `INTER` parameters together improves backup time on large systems. For example, the following `STORE` command uses both parallel device pools and interleaving to minimize backup time:

```
:STORE @. @. @.; ;SHOW=OFFLINE; &  
STORESET=( *T, *T ), ( *T, *T ); INTER
```

TurboSTORE/iX II Options

The following options are available only if the TurboSTORE/iX II product is purchased.

In addition to all of the STORE options, the following additional options are available with TurboSTORE/iX II. For more information on purchasing TurboSTORE/iX II, contact your Hewlett-Packard sales representative.

COMPRESS For file compression

Compressing Data (COMPRESS)

With TurboSTORE/iX II (and TurboSTORE/iX 7x24 True-Online Backup), backup performance can be improved through the use of host data compression. By specifying the COMPRESS parameter, the host system filters the data through a compression algorithm prior to writing it to the backup device. This operation has two effects:

- The device's throughput is increased by transferring more data than the device's maximum sustained rate. The increase is dependent upon available CPU resources.
- Fewer media changes are required by putting more data on each piece of media.

Two levels of compression are supported by the COMPRESS parameter: HIGH and LOW.

NOTE No matter whether you specify HIGH or LOW, the actual level of data compression is dependent upon the type of data being compressed. For example, program files will not compress as much as data files.

You must determine which of the compression algorithms best suits your needs. In some cases the low-compression algorithm may perform almost as well as, or better than, the high-compression algorithm.

It is advisable to try both algorithms on your data in order to evaluate which algorithm best suits your needs. Some of the factors to consider are:

- Elapsed time of backup.
- Amount of CPU time required.
- Amount of media required.

Be advised that when performing backups with the high-compression algorithm, the backup can consume much or most of the processing capabilities of the machine. The exact amount of system resources consumed depends on your system speed and multi-processor configuration. Using the low-compression algorithm demands less system resources, so it is easier for other activity to coexist with the backup.

Use the high-compression algorithm if backup time is dedicated or if the difference in the amount of media used for backup is substantial enough to warrant the complete consumption of the machine.

When compressed format files are restored, RESTORE automatically recognizes the compressed format and restores the files properly.

To compress data when storing files to backup, use the `FILE` command, then use the `COMPRESS` parameter in your `STORE` command. For example:

```
:STORE @.OPERATOR.SYS;*T;COMPRESS=HIGH
```

WARNING **If you are using a backup device that performs hardware data compression, `COMPRESS=HIGH` should NOT be used. Performing compression twice is redundant, and in some cases can lead to slower backup performance. The best throughput results occurs when using no compression or `COMPRESS=LOW` with hardware compressing devices.**

TurboSTORE/iX 7x24 True-Online Backup Options

In addition to the above options, the following options are available with TurboSTORE/iX 7x24 True-Online Backup:

- `ONLINE` After a short attach period, files can be accessed and modified during backup.
- `ONLINE=START ONLINE=END` Files can be accessed without interruption throughout the backup.
- `LOGVOLSET` Allows you to specify what volumeset should be used for online backup log files.
- Store to Disk** Instead of writing to a tape or MO device, `STORE` can write the backup to a disk file. Refer to Chapter 3, "Preparing Storage Devices," for more information.

Storing Files Online (ONLINE)

Using TurboSTORE/iX 7x24 True-Online Backup, you can back up files and applications that need to be available most or all of the time. TurboSTORE 7x24 True-Online Backup can perform two different types of online backups:

- Using the `ONLINE` option only:

The first type of online backup requires that all files being store be close for write access when the backup first starts. This is the type of online backup available on pre-MPE/iX 5.5 systems, when the TurboSTORE/iX II with Online Backup products are used (product numbers HP36388A or HP36398A). This type of backup is referred to in this manual as an *online backup*.

- Using the `ONLINE=START` or `ONLINE=END` options:

Using the new functionality of TurboSTORE/iX 7x24 True-Online Backup, you can also do an online backup of files and applications without requiring that files be closed for write access at any time during the backup. This is referred to in this manual as a *7x24*

*true-online
backup.*

- The second type of online backup, referred to as a *7x24 true-online backup*, lets you backup applications without requiring any files to be closed for write access at any time during the backup.

This section describes how to perform the online backups using just the `ONLINE` option. For information on performing 7x24 true-online backups, refer to Chapter 7.

Performing an Online Backup

At the time the online backup is started, all files must be closed for a brief period while they are attached to the *shadow logging* system. Shadow logging saves the "before images" of file blocks when the records in the file block are modified during online backup. The before images are stored on tape and are used by `RESTORE` to reconstruct the original contents of the file.

How much time is required for the files to be closed depends upon how many files are being stored. The range of time is usually from 1 to 30 minutes. Applications should be closed during this attach period. Any files open for write access during this period are **NOT** stored.

The shadow logging subsystem maintains the original state of the files as they appeared when the online backup began. If an attached file is modified, the before image of the portion (block) of the file being modified is saved in a log file. After the store process has written the attached file to the backup media, it requests the logging information from the shadow subsystem and appends the logged data, if any, to the end of the data file.

If an attached file is purged during the online backup, the system logically removes the file, but does not physically remove it until it has been stored. Therefore, the backup reflects the state of the system at the time `STORE` was initiated.

NOTE

1. The shadow logging operation can fail if there is not enough disk space. Error messages are displayed for all files that fail to store due to lack of disk space. However, the backup will **NOT** abort.
2. You must purchase TurboSTORE/iX 7x24 True-Online Backup (B5152AA) to use the `ONLINE` option. This functionality is the same as was previously available in TurboSTORE/iX II with Online Backup (HP36388A and HP36398A).
3. Any Image databases being stored with the `ONLINE` option should be closed for write access before the backup is started. If user logging is enabled for the database(s), then the current log file should be closed when the database is closed. Once the attach period has finished, the database can be opened again. A new user log file should be started at this time.

Heavily modified files should be stored first. This is accomplished by specifying these files first in the `STORE` command fileset list. By storing these files first, you reduce the system overhead and the amount of logging information written to the backup media.

When online backup begins, the shadow subsystem assigns each file in the store fileset to a log file created on the same volume where the file resides. Actual logging activity does not take place unless the file is modified after the attach phase.

TurboSTORE saves the changes made to databases and files being stored in special log files. These files can become quite large when there is a lot of activity on the system. The LOGVOLSET option allows you to specify which volume set to use for the log files. You can use this feature with the log file size summary statistics from previous backups to ensure that there is enough disk space for the online backup log files. The log volume set is specified as follows:

```
:STORE /;*tape;ONLINE;LOGVOLSET=volumeset
```

The amount of log space required is directly proportional to the amount of modifications performed against the attached file. Therefore, the more modifications made to an attached file, the more log space required. In the worst case, the maximum amount of log space required equals the sum of the space required for all of the files being stored. By storing the heavily modified files first, you reduce the amount of logging space required. Reducing the amount of logging data written to the backup media will also reduce the time required to restore the files.

As each piece of backup media is filled, logging is disabled for the files that were stored on that media.

Use the ONLINE parameter of STORE to perform an online backup. Before the STORE process is initiated, you must warn the users to close all files to ensure that the files are both physically and logically consistent. For example, the following message can be sent:

```
WARN @S;ONLINE BACKUP WILL BEGIN IN TEN MINUTES.  
WARN @S;PLEASE CLOSE YOUR FILES.
```

Begin the online backup by specifying the ONLINE parameter with the STORE command. For example, the following illustrates the STORE command for an online backup.

```
:STORE @.@@;*T;DIRECTORY;ONVS=MPEXL_SYSTEM_VOLUME_SET,NV1;ONLINE
```

When online backup begins, all files not currently open for write are attached to a log handler known as the shadow subsystem. Once the attach phase is complete, STORE issues the following message to the console.

```
FILES LOCKED BY ONLINE STORE ARE NOW FREE FOR READ/WRITE/PURGE
```

This message is also written to \$STDLIST if STORE is being run from a session. After you receive the above message, the users are allowed read, write and purge access to the attached files. Send your users a message telling them they may resume their applications. For example:

```
WARN @S;ONLINE BACKUP HAS BEGUN.  
YOU MAY NOW RESUME YOUR APPLICATIONS.
```

The NOTIFY option (described earlier in this chapter) can be used to automate this task.

If any files were still open at the time you started online backup, you will receive the following message:

```
! NOT STORED: FILE OPEN FOR WRITE AT ATTACH
```

where the exclamation point (!) is replaced by the actual filename(s).

During the RESTORE process, the log file is used to reconstruct the original contents of the files; that is, the way they appeared at the moment they were attached during online backup.

NOTE The ONLINE option cannot be specified with the SPLITVS or TRANSPORT options.

Specifying Where to Place Log Files (LOGVOLSET)

TurboSTORE/iX 7x24 True-Online Backup saves the changes made to databases and files being stored in special log files. These files can become quite large when there is a lot of activity on the system. The LOGVOLSET option allows you to specify which volume set to use for the log files. You can utilize this feature with the log file size summary statistics from previous backups to ensure that there is enough disk space for the online backup log files. The log volume set is specified as follows:

```
:STORE /;*tape;ONLINE;LOGVOLSET=volumeset
```

LOGVOLSET can be used during an online backup (ONLINE) or a 7x24 true-online backup (ONLINE=START or ONLINE=END). It has no effect when specified for a non-online backup.

7 TurboSTORE/iX 7x24 True-Online Backup

This chapter describes how to use TurboSTORE/iX 7x24 True-Online Backup. It includes a description of all the options provided for performing 7x24 true-online backups, as well as discussions about TurboSTORE's interactions with the system during these backups.

The sections in this chapter include:

- Using TurboSTORE/iX 7x24 True-Online Backup
- Consistent Backups
- How Shadow Logging Works
- Logging File Attribute Modifications
- TurboSTORE/iX 7x24 True-Online Backup Options
- Restoring TurboSTORE/iX 7x24 True-Online Backups

Using TurboSTORE/iX 7x24 True-Online Backup

When using TurboSTORE/iX 7x24 True-Online Backup, you can perform file, application, and system backups without closing files or disrupting users. Some guidelines for performing these backups are discussed below.

TurboSTORE 7x24 True-Online Backup can perform two types of online backups:

- The first type is referred to as an *online backup*, which requires that all files being stored be closed for write access when the backup first starts. This is the type of online backup formally available with the TurboSTORE/iX II with Online Backup product.

The online backup refers to a backup created with only the `ONLINE` option of TurboSTORE/iX 7x24 True-Online Backup.

- The second type is referred to as a *7x24 true-online backup*, which lets you backup applications without requiring any files to be closed for write access at any time during the backup.

The 7x24 true-online backup refers to a backup created with the `ONLINE=START` or the `ONLINE=END` option of TurboSTORE/iX 7x24 True-Online Backup.

The Sync Point

When performing a backup, you want to know at what point in time the backup occurred. This is helpful when performing a recovery so you know that all files were restored to a state at the same time. This is the point called the *sync point*, where all files are synchronized. It can occur at different times, depending on the type of backup performed.

You can schedule when the sync point occurs (either at the start or the end) by scheduling when the backup job is launched. Be aware, however, that a certain set up time is required for TurboSTORE to create the list of files to be stored. Depending on the number of files being stored, the sync point may happen soon after the backup is launched, or may occur 5-15 minutes later. To better control when the sync point occurs, a fixed time can be specified on the `STORE` command line. See "Creating a 7x24 True-Online Backup (ONLINE)" later in this chapter for more details.

The `NOTIFY` option can also be used to launch a job stream once the sync point has occurred. See Chapter 6, "Notification of File Availability (NOTIFY)".

WARNING All 7x24 true-online backups created with the sync point at the beginning can be restored on any MPE/iX system. However, backups created with the sync point at the end can only be restored or verified on MPE/iX Release 5.5 or later. If you know at the time you perform the backup that a tape must be restored onto an earlier release, create the tape with the sync point at the beginning.

Sync Point at the Beginning of Backup

The sync point is always at the beginning of an online backup when only the `ONLINE` option is used. A 7x24 true-online backup can also have the sync point at the beginning by using the `ONLINE=START` option.

The main difference between these two types of backups is that an online backup (`ONLINE`) requires all files to be closed at the sync point, while a 7x24 true-online backup (`ONLINE=START`) allows users to access and modify their files, uninterrupted, throughout the sync point and backup.

Why Use the Sync Point at the Beginning of a Backup?

Using the sync point at the beginning of a backup has the following advantages over using it at the end:

- Synchronizing at the start of a backup gives you better control of when the sync occurs.
- It spreads the log data throughout the tape and therefore one is less vulnerable to media errors.
- It allows for faster recovery of partial restores, since sync at end requires `RESTORE` to read the log files from the end of the last piece of media.
- It allows for migration of files to earlier releases of TurboSTORE/iX, whereas sync at end does not.

For example, scheduling the sync point at the beginning of the backup is useful when your system is heavily used during the day, making performing backups then impractical. The backup can be started at the end of the busy day shift, and run during the night. Since no application shutdown is required, users, perhaps from other time zones, can continue to access their files and applications uninterrupted.

Sync Point at the End of Backup

TurboSTORE/iX 7x24 True-Online Backup can also schedule the sync point at the end of the backup using the `ONLINE=END` option. Using the options described in "Creating a 7x24 True-Online Backup (`ONLINE`)", you can schedule the sync point to occur at a fixed time, or at the end of the backup. However, if you specify the sync point to occur at a certain time, it does not occur until all files have been stored.

Why Use the Sync Point at the End of a Backup?

Using the sync point at the end of a backup has the following advantages over using it at the beginning:

- Synchronizing at the end of the backup creates a logically consistent copy of your data at the time of the end of the backup.
- It gives you the most recent copy of your data.

The backup media can be off the system and stored in a data safe before running a "risky" application. If the "risky" application fails, you can recover your data at the point just prior to starting the application.

- All file changes occurring during the backup are logged and stored at the end of the

backup media.

For example, scheduling the sync point at the end of the backup can be useful if you start your backup in the middle of the day (at lunch time, for example). The backup can then run throughout the afternoon shift, and be finished before your disk-intensive batch jobs are run at night. Synchronizing at the end would then allow you to save the state of your system at the end of the working day.

If your batch jobs do not run successfully, you can restore from that backup without losing a day's work. If this is important to business operations, you may want to consider backup with the sync point at the end.

Consistent Backups

Even though database files generally constitute the largest part of any backup, TurboSTORE/iX 7x24 True-Online Backup also backs up system and non-database files. Since TurboSTORE allows changes to files during a backup, it establishes a close tie to the file system to ensure physical and logical consistency.

Physical Consistency

Physical consistency is ensured by providing a snapshot of the fileset at a specific point in time, specifically, at the sync point. This is adequate for files that are not "related" to other files, and have no dependencies on changes to other files. This includes system files, some flat files and any other files not associated with a database.

For KSAM XL files and files enabled for transaction management, TurboSTORE guarantees that no partial transactions are included in the copy of the file stored on the tape. This ensures that a consistent copy of the file is stored.

If Native Mode Message files are open during the sync point, TurboSTORE follows the necessary steps to successfully imitate closing the file to ensure data integrity, physical consistency, and recoverability.

User-mapped files are managed the same way they were with as TurboSTORE/iX II with Online Backup. Whenever changes are made to a user-mapped file being stored by a 7x24 true-online backup, the entire file is copied to TurboSTORE's log files. However, if a user-mapped file is left open for write access across the sync point, then TurboSTORE captures the last known state of the file, which would either be at the last FOPEN, FCLOSE, or FCONTROL 6. If TurboSTORE does not capture any of these events, the file is stored in whatever state it was in when TurboSTORE opened the file, which may be inconsistent. A warning message is printed for all user-mapped files open across the sync point.

NOTE Compatibility mode (CM) files, namely Circular (CIR) files, and Relative I/O (RIO) files have internal buffer structures that cannot be logged by TurboSTORE. If these files are open for write access during the sync point, TurboSTORE captures their last known state, which is either since the last FCLOSE or the last FCONTROL 6. A warning is issued for all CM files that were open across the sync point.

Logical Consistency

In addition to physical consistency, TurboSTORE/iX 7x24 True-Online Backup also guarantees logical consistency for databases. This is done by setting the sync point to occur when all open transactions against the databases have completed. No new transactions can start until the open transactions are complete, which is called *quiescing* the database.

Although any type of database can be stored by TurboSTORE, only TurboIMAGE and ALLBASE/SQL databases are automatically quiesced using special database intrinsics.

NOTE Quiescing an ALLBASE/SQL database and storing the ALLBASE/SQL database by TurboSTORE/iX 7x24 True-Online Backup when the DBEnvironment is in normal use, has the same effect as using the `STOREONLINE` command of SQLUtil. For a 7x24 true-online backup of ALLBASE/SQL to be of use for rollforward recovery, use the same procedures you would use when doing the `STOREONLINE` commands.

If you need to synchronize any other file types (e.g., CM file types such as CIR files and RIO files) and/or any third-party databases that are not automatically synchronized, you can manually quiesce these before beginning the backup, or during the backup by specifying the `ASK` parameter with the `ONLINE` option. `ASK` makes TurboSTORE wait at the sync point while you finish these operations. Refer to "Prompting During Backup (ASK)" later in this chapter for a description of the `ASK` parameter.

Once all files and databases are quiesced, TurboSTORE establishes the logically consistent synchronization point. It then unquiesces all TurboIMAGE and ALLBASE/SQL databases.

How Shadow Logging Works

TurboSTORE/iX 7x24 True-Online Backup uses *shadow logging* to intercept any changes made to files while they are being stored.

When the Sync Point is at the Beginning

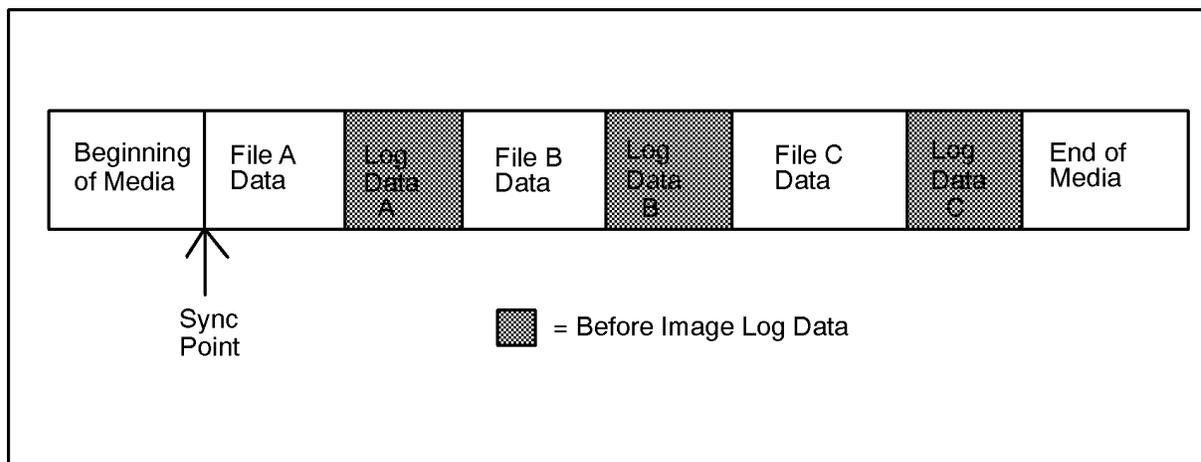
When the sync point occurs at the beginning of the backup (`ONLINE=START`), shadow logging guarantees that TurboSTORE captures the state of the fileset as it is right before the files begin to be written out. When a file is opened for write access, any changes made to the file are monitored by shadow logging. When a modification occurs, the part of the file being changed is copied. This "before image" is added to internal shadow log files. Once the data has been copied, the modification can proceed. The before image logs for all files are maintained until the information is requested by TurboSTORE. TurboSTORE requests a file's log data once it has finished storing the file. The log data is written immediately following the file data. The before image log data is maintained until the media that the file resides on has been completed. At reel switch time, or at the end of the backup (whichever comes first), the before image log data for all files on that media is released.

By default, the before image log files are created on the same volume set as the files being stored. You can create the log files on a specific volume set by specifying the volume set name using the `LOGVOLSET` option, as explained in "Performing Online Backup" in Chapter 6, "STORE and TurboSTORE Command Options."

`RESTORE` restores a file by first reading the file data that had been written to the media. The log data is then read in and replaces any sections of the file that were already restored. Once a file's before image logs are applied, the file is restored to its state at the time of the sync point.

Figure 7-1 shows what a backup tape looks like if the sync point occurs at the beginning of the backup (`ONLINE=START`).

Figure 7-1. Backup Created with Sync Point at the Beginning



When the Sync Point is at the End

When the sync point occurs at the end of the backup (`ONLINE=END`), shadow logging captures the state of the files as they are at the end of the backup. In this case, instead of copying the before images of modifications to log files, shadow logging uses a bitmap to keep track of the modified pages. Once the backup is completed, the bitmap is used to collect the modified "after image" pages from the system. This data is then written out to the end of the media. For parallel backups, after image log data is written at the end of each parallel set. For more information on parallel backups, see Chapter 3, "Preparing Storage Devices".

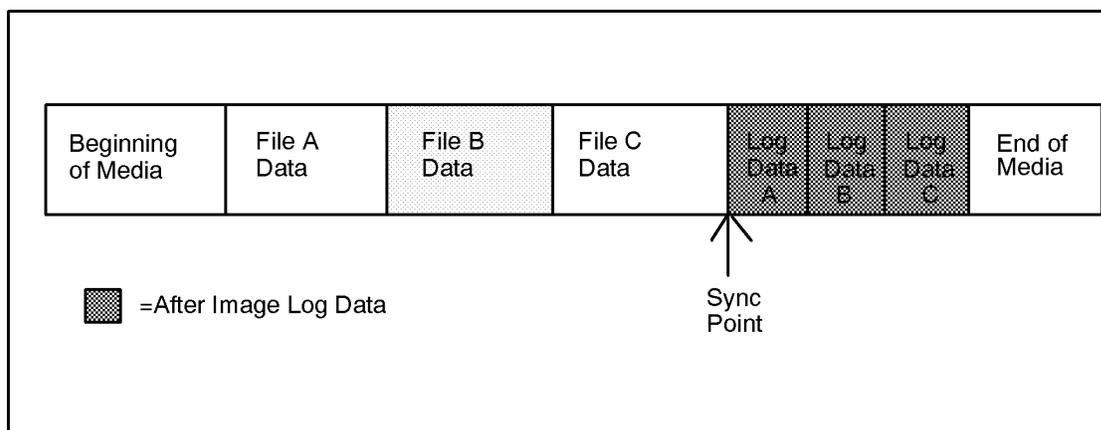
During the backup when `ONLINE=END` is specified, logging is turned on for a file when TurboSTORE begins to write it out. The logging stays on for a file until the sync point occurs at the end of the backup.

When a file with after image log data is restored, the initial file data is read in first. Once the after image log data is read in, it is applied to the file data, reapplying all changes made to the files during the backup. This restores all files to their state at the end of the backup.

However, since the log data is at the end of the backup, files with after image (AI) data are not available until the Restore completes and the after image log data is applied.

Figure 7-2 shows what a backup tape looks like if the sync point occurs at the end of the backup.

Figure 7-2. Backup Created with Sync Point at the End



Logging File Attribute Modifications

In addition to keeping track of changes to a file's data during a 7x24 true-online backup, changes to a file's attributes, such as filename, file limit, and file code, are also captured to ensure a consistent backup.

File Creations and Purges

When the sync point is at the end of the backup, file creations and file purges that occur during the backup are captured by shadow logging. TurboSTORE/iX 7x24 True-Online Backup also handles equivalent POSIX commands for purging files, such as `UNLINK()`, `FUNLINK()`, and other MPE/iX commands that purge files, like `PURGEACCT` and `PURGEGROUP`.

Other File Attributes

If `RENAME`ing a file causes it to qualify for a 7x24 true-online backup, the file is included in the backup. Conversely, if a file is renamed to not qualify for a 7x24 true-online backup's fileset list, it is removed from the backup. These actions only happen if the sync point is at the end of the backup.

TurboSTORE/iX 7x24 True-Online Backup Options

This section discusses the new options available for creating TurboSTORE/iX 7x24 True-Online Backups.

Creating a TurboSTORE/iX 7x24 True-Online Backup (ONLINE)

A TurboSTORE/iX 7x24 True-Online Backup is performed only when either the `START` or `END` parameter is included with the `ONLINE` option.

If no parameters are specified for `ONLINE`, the sync point automatically occurs at the beginning of the backup, and an online backup is performed where all files must be closed for write access at that time.

Up to 15 TurboSTORE/iX 7x24 True-Online Backups can be running at once. Also, a file can be stored by up to 15 7x24 true-online backups.

Initiating a TurboSTORE/iX 7x24 True-Online Backup (START and END)

The initiation of a 7x24 true-online backup with a logically consistent sync point is controlled by the `START` or `END` parameter with the `ONLINE` option. Backups created with the `START` parameter have the sync point at the beginning of the backup. Backups created with the `END` parameter have the sync point at the end of the backup.

Specifying Time

In addition to the `START` or `END` parameter, you can specify a time for the sync point to occur. The time is specified in 24-hour format (hh:mm[:ss]). If a time is specified with `END`, the sync point occurs at the time specified or once all files have been stored, whichever event occurs last. If a time is specified with `START`, the sync point occurs at the time specified, or once all files are prepared for the backup, whichever event occurs last.

For example, you can initiate a backup with the following command:

```
:STORE /;*tape;ONLINE=END,17:30
```

If the backup completed at 17:00, TurboSTORE would continue to log changes until 17:30. It would then synchronize the databases and write the log data to the backup media. If, on the other hand, the backup did not complete until 18:00, TurboSTORE would delay the sync point until 18:00. See the "Synchronizing Multiple Applications" section later on, for more information on synchronizing all applications.

Prompting During Backup (ASK)

If you want to be prompted during the backup to manually quiesce any other file types and/or third-party databases that are not automatically synchronized, you can use the `ASK` parameter. `ASK` makes TurboSTORE wait for you to finish these operations.

NOTE If you are running the backup unattended, and you want to synchronize these file types and/or databases, manually quiesce these before beginning the backup.

If the following backup is initiated:

```
:STORE /;*tape;ONLINE=END,17:30,ASK
```

TurboSTORE asks:

```
17:30/#S21/59/ARE ALL FILES AND DATABASES QUIESCED FOR ONLINE  
BACKUP (Y/N)?
```

When all of your applications are quiesced, you should reply:

```
:REPLY 59,Y
```

If you reply with the following:

```
:REPLY 59,N
```

TurboSTORE asks:

```
17:34/#S21/59/DO YOU WANT TO ABORT THE BACKUP(YES/NO)?
```

A reply of YES aborts STORE and the backup is incomplete and invalid. A reply of NO causes TurboSTORE to ask the original question:

```
17:34/#S21/59/ARE ALL FILES AND DATABASES QUIESCED FOR ONLINE  
BACKUP (Y/N)?
```

When TurboSTORE receives a reply of YES to this question, it quiesces all TurboIMAGE and ALLBASE/SQL databases, turns on before image logging, unquiesces the databases, tells you to unquiesce all your applications, and completes the backup. You are notified on the console:

```
17:35/#S21/59/ALL OF YOUR APPLICATIONS CAN BE UNQUIESCED NOW.
```

You can then perform any actions needed to unquiesce your applications.

Synchronizing Multiple Applications

You can synchronize multiple applications in several ways. Which method you use depends on the situation and how closely you want the applications to be synchronized. When you perform multiple stores per application, they can only be approximately synchronized.

- Store all applications with a single STORE to ensure applications are synchronized. All databases are quiesced at a single point in time.

This is the preferred method, since it synchronizes the applications. The only possible disadvantage is that restoring one of the applications may take longer.

- Use the END and ASK parameters to approximately synchronize applications.

When you use ONLINE=END,ASK, all jobs are backed up and the operator request for each backup appears on the console. Wait for all requests for the backup jobs to be displayed. Then reply to each request to complete all the backups, which causes the backups to be approximately synchronized.

If a job fails, ASK would not show a request for that job. You could first do a `RECALL` to see if all the requests are there. If the job is missing, you could go fix the job, manually rerun it until it gets to ASK, and then complete all the backups by replying to all the ASK operator requests.

- Use the `START` or `END` parameter with the time parameter.

When you use `ONLINE=START` with the time parameter, set the time far enough in the future to be sure the setup time for each backup has completed. When you use `ONLINE=END` with the time parameter, set the time far enough in the future to be sure all backups are complete.

Using `START` or `END` with the time parameter is a good method for synchronizing all stores sequentially, causing them to have the same approximate sync points.

- Use the `END`, `ASK`, and time parameters.

If you think there might be a job failure, this is the best method to use. Set the time far enough in the future to be sure all backups are complete. This method puts all the ASK requests on the console at the specified time, instead of displaying each job request as it completes.

Roll-forward Recovery of TurboIMAGE/XL Databases (DBRECOV)

DBRECOV performs rollforward recovery of TurboIMAGE/XL databases stored using the new TurboSTORE/iX 7x24 True-Online Backup (with either the `START` or `END` option). Databases enabled for roll-back recovery are not affected by using the 7x24 true-online backup.

To perform rollforward recovery of your database stored using 7x24 true-online backup, the following requirements must be adhered to:

- Your system must be up on MPE/iX Release 5.5 or later.
- `LOGID` must include the `AUTO` option. If it was not used in the `GETLOG` command, use the `ALTLOG` command to include the `AUTO` option.

```
:GETLOG imageid;log=image001,DISC;pass=wontell;auto
```

or

```
:ALTLOG noautoid;auto
```

- User Logging Process must remain active when storing the database even when the database is not open. That is, `LOG logid,STOP` must not be issued before storing the database. This is because the logging information which is dynamic in nature, is incorporated in the root file when the database is stored. This dynamic logging information can only be obtained when the log process is active. It is used later when performing rollforward recovery.
- After the backup is completed, purge all logfiles preceding the one in use when the backup was initiated. It is essential that you retain the one used when backup was initiated and the ones following that. These are the logfiles that will be needed later to perform rollforward recovery.

For further information on DBRECOV, see the MPE/iX Release 5.5 (C.55.00) *Communicator* article, "IMAGE/SQL\with TurboIMAGE/XL Enhancements."

Storing Part of a Database (PARTIALDB)

When doing a 7x24 true-online backup, by default, TurboSTORE does not allow you to specify an incomplete TurboIMAGE or ALLBASE/SQL database to be stored. If you want to override this behavior, specify the PARTIALDB option. Any databases that are partially stored are **NOT** quiesced at the sync point.

NOTE Database corruption may result if not all database files are stored during a 7x24 true-online backup. Be sure that you only want to store certain database files before overriding the default behavior with PARTIALDB.

Refer to the "Storing Database Files Using PARTIALDB or FULLDB" section in Chapter 6 for more information on using this option.

Creating a Disk Directory File (STOREDIRECTORY and NOSTOREDIRECTORY)

When creating a 7x24 true-online backup with the sync point at the end of the backup, the STOREDIRECTORY option is enabled by default. To override this behavior, specify the NOSTOREDIRECTORY option. It is strongly recommended that you use the STOREDIRECTORY option when creating true-online backups with the sync point at the end of the backup. Doing so greatly increases recovery time when the backup needs to be restored.

Refer to the "Creating a Disk Directory File (STOREDIRECTORY and NOSTOREDIRECTORY)" section in Chapter 6 for more detailed information on using the STOREDIRECTORY option.

Restoring TurboSTORE/iX 7x24 True-Online Backups

This section discusses the options available for restoring backups created by TurboSTORE/iX 7x24 True-Online Backup.

When restoring a backup that was created with the sync point at the beginning of the backup, RESTORE performs the following tasks for each file restored:

1. Restores the initial file data.
2. Applies any before image log data that exists on the backup.

When a 7x24 true-online backup with the sync point at the end of the backup is restored, RESTORE performs the following tasks:

1. Restores the initially stored files.
2. Restores all new files that were added to the fileset during the backup due to renames or builds.
3. Applies any changes made to files after they were written to the media (the after image logs).
4. Does not restore all files that were deleted from the fileset during the backup.

If RESTORE is unable to read the media directory information from a media directory disk file, it asks you to mount the last piece of media from the backup. Once it has read the media directory, it prompts you to mount the necessary piece(s) of media for the recovery process.

NOTE	Since the after image log data is stored at the end of the backup, files that need to have after image log data applied is unavailable for read OR write access, until the after image log data is applied at the end of the backup. Once RESTORE reads and applies the after image log data for a file, the file can be accessed again.
-------------	--

Restoring with a Disk Directory File (STOREDIRECTORY)

When restoring files, RESTORE first reads in the Store labels from the devices specified either in the restorefile parameter, or in the RESTORESET option. It then opens the appropriate media directory disk file, if one exists. This allows RESTORE to more efficiently restore files. The greatest advantage occurs when selectively restoring files from many tapes, or if the backup was created with the sync point at the end.

Selective restores can be optimized further by specifying the disk directory filename with the STOREDIR[ECTORY] option. If you specified a name for the backup at Store time with the STOREDIRECTORY option, that name can also be specified here. Alternately, you can specify the full HFS directory disk file created by STORE.

Using the STOREDIR option causes RESTORE to read the directory on disk and prompts you to mount the correct media for the restore, shown as follows:

```
17:30/#S20/40/MOUNT MEDIA 2 OF SET 1 FOR RESTORE ON LDEV 7
```

For example, to restore the backup created on June 22, 1994 at 1:30pm by Pin 30, issue the following command:

```
:RESTORE *tape;JONNYS.BIG.FILE;&
```

```
:STOREDIR=/SYS/HPSTORE/store_dirs/store_19940622_1330_pin30_wed
```

To restore the full backup from Friday named FRIFULL, issue the command:

```
:RESTORE *tape;//STOREDIRECTORY=FRIFULL.HPSTORE.SYS
```


8 Handling Media Requests

This chapter describes the processes for handling the backup media. The subjects include:

- Preparing the Backup Device
- Checking the Console
- Replying to a Media Request
- Pending Media Requests
- Denying a Media Request
- Using `REPLY` and `RECALL`
- Automatically Assigning Devices

NOTE Instructions for using backup devices in this chapter are general procedures that apply to most devices. Your devices may operate differently. Refer to the documentation for each device for exact instructions and additional information.

Unless your backup devices are configured to be automatically allocated, when a user (or a user's program) needs to use a backup device, you see a media request at the console. For example:

```
?16:04/#S37/23/LDEV# FOR "T" ON MEDIA (NUM)?
```

The media request asks you to assign an LDEV number to the file named in the request. This process is called allocating the device. The request above asks you to assign a logical device to the file `T`.

It tells you the time it arrived (16:04 in the example above), the job or session that sent it (#S37 in the example above), and the process identification number (called the PIN) of the program or command that sent the request (23 in the example).

Depending upon your system configuration, the system may or may not answer media requests automatically. Refer to "Automatically Assigning Devices" later in this chapter. If it does, you need only to prepare the backup device for use. If the system does not answer media requests automatically, you must prepare the backup device and reply as explained in "Replying to a Media Request," below.

Preparing the Backup Device

In general, a backup device cannot be shared. Only one person at a time can use it. Often, users cannot access media drives without your assistance because media drives may be in a restricted computer room or you may be required to monitor media drive use.

Users only have direct access to media drives when you provide them physical access and configure media drives so that they are automatically assigned to users. See "Automatically Assigning Devices," later in this chapter.

When a user (or a program) requests the use of a media drive and the device is available, you mount the media so that the storage or Restore process can begin. You mount the media by preparing it, loading it onto the media drive, and placing the media drive online.

To prepare a backup device for storing or restoring files, use the media from the user you are helping (or obtain it from your media library) and mount it on the device.

When the media is ready, you see a `VOLUME MOUNTED` message on the console. For example:

```
16:05/31/VOL (Unlabeled) mounted on LDEV#8
```

The `VOLUME MOUNTED` message indicates the logical device (LDEV) number of the media drive. In the example above, the LDEV number is 8. Note the LDEV number. You use it to assign the device to `STORE` or `RESTORE` when you respond to the media request.

Checking the Console

During periods you are away from the console, other messages and requests can cause media requests to scroll off the screen. When you return to the console, some messages and requests will no longer be in view. The system keeps track of media requests until you answer them. Use the `RECALL` command to check for outstanding media requests. For example:

```
:RECALL
```

```
THE FOLLOWING REPLIES ARE PENDING:  
?16:04/#S37/23/LDEV# FOR "T" ON MEDIA (NUM)?
```

Replying to a Media Request

Check that the request comes from the user who gave you the media. First, locate the user's job or session number in the media request. There can be different PINs for each request. In the following example, the session number is #S37.

```
$$?16:04/#S37/23/LDEV# FOR "T" ON MEDIA (NUM)?
```

Issue a SHOWJOB command to determine from which user the request came. For example:

```
:SHOWJOB #S37
```

The system describes the job or session. For example:

Figure 8-1. Job Session Description

```
JOBNUM   STATE   INPRI   JIN    JLIST   INTRODUCED  JOB NAME
#S37     EXEC                36     36      MON8:46A
        WANDA.PERSONEL
```

```
JOBFENCE= 5; JLIMIT= 10; SLIMIT= 40
```

Check that the job name matches the user's job name. If it does, use the REPLY command to assign the backup device to the user.

To reply to a media request, you need to know the STORE or RESTORE user's PIN and the media drive's LDEV number. For example, the PIN for the following media request is 23.

```
?16:04/#S37/23/LDEV# FOR "T" ON MEDIA (NUM)?
```

The backup device's LDEV number appears in the message you receive when you mount a media on the backup device. In the following message, the LDEV number is 8.

```
16:05/31/VOL (Unlabeled) mounted on LDEV#8
```

To answer a media request, use the REPLY command to assign the LDEV number to the PIN. For example, the following REPLY command assigns LDEV 8 to PIN 23.

```
:REPLY 23,8
```

Pending Media Requests

If several media requests appear at once, you can choose to respond to any or all of them, in any order. When you choose not to respond to a media request, it remains pending until you respond or deny it. You can use the `RECALL` command at any time to list pending media requests.

Denying a Media Request

If the backup device is unavailable, or for another appropriate reason, you choose not to let a user or users access a backup device, you should deny media requests when they appear. To deny a media request, use the `REPLY` command to assign the `LDEV` number 0 to the user's PIN. So, for example, to deny a media request you received from the user with PIN 23, enter:

```
:REPLY 23,0
```

The user receives a message reporting that the device (the media drive) is unavailable. You may want to send the user another message using `TELL` explaining the situation in detail and estimating when the device might be available.

Using `REPLY` and `RECALL`

During the time you are running a program on the console, you are unable to get a colon prompt on the console. However, you can still review and reply to media requests. At the console, press the `CTRL` and `A` keys simultaneously to display the equal-sign (=) prompt. In response to the equal-sign prompt, you can issue either the `RECALL` or the `REPLY` command, described in the preceding sections. This feature is useful, for example, when running a program such as `STORE` which you expect will take a long time to complete.

Automatically Assigning Devices

You may choose to configure backup devices so that they are automatically assigned to users. This is also referred to as *auto reply*. When devices are configured for automatic assignment, the system automatically answers media requests for you.

Devices configured for automatic assignment must have the following characteristics:

- The device mode must be R (autoreply).
- The device must not be job or data accepting.
- The device must not be assigned to another user.
- The device specification must be unique. You must use the device specification (logical device number or unique device class) in a `FILE` command describing the device.
- The user must request an unlabeled media.

Refer to *System Startup, Configuration, and Shutdown Reference Manual* for instructions for configuring devices.

In order for the system to automatically assign a correctly configured device, the user (or you, if you are storing or restoring files yourself) must name the device and its LDEV number (or device class, if the device class is unique) in a `FILE` command preceding the `STORE` command, and the device must not be assigned to any other user. If a device cannot be automatically assigned, you must assign the device in the standard manner.

Instructions for using the `FILE` and `STORE` commands can be found in Chapter 5, "Storing Files."

9 When STORE is Complete

When you see the `END OF PROGRAM` message, `STORE` is complete. To be absolutely sure that important files have been stored correctly, you will probably want to do the procedures in following sections:

- Checking File Listings for Errors
- Verifying That `STORE` Files Are Error Free
- Using `NODECOMPRESS`
- Retaining the `STORE` Tapes
- Allowing Users Back on the System

Checking File Listings for Errors

As STORE copies files to backup, it will list them on the \$STDLIST in the SYSLIST file, provided you use the SHOW parameter. For example, if you enter

```
:STORE T@.SREXP.SYS;*T;SHOW
```

The following might be listed on the console:

Figure 9-1. Listing Files Copied to Tape

```
STORE/RESTORE VERSION X.51.00 (C) 1986 HEWLETT-PACKARD CO.
```

```
STORE a[mn]@.pub.sys;*n;SHOW
```

```
MON, JUL 10, 1995, 5:50 PM
```

FILENAME	GROUP	ACCOUNT	VOLUME RESTRICTIONS	SECTORS	CODE	MEDIA
AMCAT000	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	112	MGCAT 1
AMMESCAT	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	96	1
AMSYSDIC	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	1744	1
AMTCX	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	384	NMPRG 1
AMTRX	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	1840	NMPRG 1
ANSTART	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	16	1
ANSTAT	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	16	1
ANSTOP	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	16	1
ANUTIL	.PUB	.SYS	MPEXL_SYSTEM_VOLUME_SET	:S	288	NMPRG 1

```
FILES STORED : 9
```

```
TOTAL MEDIA WRITTEN : 1
```

When STORE finishes, you see the total number of files stored and, if there is an error, the number not stored. Use the listing to verify that the correct files were indeed stored and that none were omitted.

Verifying That STORE Files Are Error Free Using VSTORE

After using the STORE program to back up files, you should check to be sure that no errors occurred during the backup by using VSTORE. VSTORE verifies that the data written to backup is valid.

The first place to check is the STORE listing. At the end of the listing, STORE prints a summary of how many files, HFS directories, symbolic links, and MPE directories were stored. It will also display the number of files that were not stored, and the number of files that failed to store. The "files not stored" count includes all files that were specified on the STORE command line, or in an indirect file, that failed to qualify for the backup. Reasons for this could include:

- File is open for write access when an online backup was performed.
- File does not meet a DATES or FCODE specification
- User did not have appropriate security requirements to store file

The "files failed to store" count indicates a more serious problem during the backup. These files qualified for the backup, but an error occurred while actually storing the file to the backup media. The STORE listing should be examined to determine which files failed to store, and why.

The user interface for the VSTORE command is similar to that for the RESTORE command which is described in Chapter 11, "Restoring Files." Command syntax for VSTORE is shown in Appendix B, "Command Syntax."

To verify one or more files on a STORE backup, mount the backup on your backup device. The write ring should be removed as a precaution.

As with the STORE and RESTORE commands, you may optionally precede the VSTORE command with a FILE command assigning a file name to the backup drive. If you do not assign a filename, a default name is used.

For example, to verify the files on a backup, enter:

```
:FILE T;DEV=TAPE
```

Now issue a VSTORE command that backreferences the backup file and includes the file or files you want to verify. For example:

```
:VSTORE *T;@.@@
```

If you have omitted the FILE command, you would enter:

```
:VSTORE ;@.@@
```

In addition to the backup filename and the names of the files to be verified, four parameters are available: SHOW, ONERR[OR], LOCAL, and DIRECTORY. The following is a brief description of each of these parameters:

- The SHOW parameter allows you to list the names of each file successfully verified. (If you do not specify this parameter, the total number of files verified is displayed.) Additional options used with the SHOW parameter include short and long form descriptions of files, creation, last access, last modification date, security status, and an

offline printing option.

- The ONERR[OR] parameter informs VSTORE what to do if a media error is encountered while verifying a file. Two options are available. If ONERR=SKIP is specified, then VSTORE will continue to verify files after an error is encountered. It will print an error message for any file that it was unable to verify. This is the default. Specifying ONERR=QUIT causes VSTORE to terminate upon getting a media error. The rest of the files in the backup will not be verified, and an error message for each file will be issued.
- The DIRECTORY parameter allows for the file system directory on the backup to be verified. Use of this parameter requires SM or OP capability.

The following are examples of the above parameters:

```
:VSTORE *T;@. @. @;SHOW
:VSTORE *T;FILE1.MANAGER.SYS;SHOW;ONERR=SKIP
:VSTORE *T;@. @. @;DIRECTORY
```

NOTE The VSTORE command is not valid for a STORE backup created with the TRANSPORT parameter.

For a full discussion of all VSTORE command options, refer to the *MPE/iX Commands Reference Manual Volumes 1 and 2 (32650-60238)* or refer to Chapter 11, "Restoring Files." Implementation of the VSTORE options is the same as implementation of the RESTORE options, except for the ones described above. The VSTORE options are:

- SHOW
- ONERR
- DIRECTORY
- PROGRESS
- COPYACD, NOACD
- TREE, NOTREE
- NODECOMPRESS
- STOREDIRECTORY
- PARTIALDB, TurboSTORE only
- RESTORESET, TurboSTORE only
- MOSET, TurboSTORE only
- NAME, TurboSTORE only

Using NODECOMPRESS

Normally, when `VSTORE` verifies a compressed store media, it reads in each record and decompresses it. This results in a thorough verification of the media by checking both the integrity of the data on the media as well as the correctness of the compression algorithm. If the `NODECOMPRESS` option is specified, `VSTORE` will only read the records from the media. It will not decompress them. This verifies only that there are no media errors on the media. It also speeds up the `VSTORE` process, since it does not have to perform the decompression.

This option should be used when the time to verify compressed store media is a prime concern. In most cases, full verification should be done to ensure that the data was compressed and written without error.

The `NODECOMPRESS` option can only be used with `VSTORE`, for example:

```
:FILE T;DEV=TAPE
:VSTORE *T;@. @. @;SHOW;NODECOMPRESS
```

Retaining the STORE Tapes

Unload each backup and check for correct labeling. If you used the `SHOW=OFFLINE` option of the `STORE` command, the system has printed a list of the files you stored. Retrieve the report from your printer and either file it with your records or attach it directly to one or more backup reels. Place each backup in your backup library.

Allowing Users Back on the System

If you prevented users from accessing the system during your backup, as discussed in Chapter 4, "Preparing the System," you need to allow users to access the system again.

To allow users access to the system again, perform the following:

1. Reset the job and sessions limits to their original values. Enter:

```
:LIMIT nn,nn
```

2. To check that you've used the correct numbers, enter:

```
:SHOWJOB STATUS
```

The last line should report the correct job limit (JLIMIT) and session limit (SLIMIT) for your computer. If the values are incorrect enter the LIMIT command again with the corrected values.

3. Reset the system jobfence to its original value by entering:

```
:JOBFENCE nn
```

4. To check that the jobfence is correct, enter:

```
:SHOWJOB STATUS
```

If the jobfence value is incorrect, enter the JOBFENCE command with the correct value.

5. Restart any suspended jobs. To determine whether any jobs had been temporarily stopped, enter:

```
:SHOWJOB
```

The system lists your session and any suspended jobs in the following way:

Figure 9-2. Listing Suspended Jobs

```
JOBNUM  STATE  IPRI  JIN    JLIST    INTRODUCED  JOB NAME
#S185   EXEC    20    20     FRI  1:03P  OPERATOR.SYS
#J17    SUSP    10S   PP     FRI  5:02P  AJOB,MRS.T
2 JOBS:
  0 INTRO; 0 SCHEDULED
  0 WAIT; INCL 0 DEFERRED
  1 EXEC; INCL 1 SESSIONS
  1 SUSP
```

```
JOBFENCE= nn; JLIMIT= nn; SLIMIT= nn
```

The example shows one suspended job (#J17). To restart that job, enter:

```
:RESUMEJOB #J17
```

Repeat the RESUMEJOB command for each job listed as SUSP. When you have done this for each suspended job, check that all have been restarted by entering:

```
:SHOWJOB
```

The list should look nearly identical, except that jobs that were suspended before

should now be listed as executing or EXEC:

Figure 9-3. Listing Resumed Jobs

JOBNUM	STATE	IPRI	JIN	JLIST	INTRODUCED	JOB NAME
#S185	EXEC		20	20	FRI 1:03P	OPERATOR.SYS
#J17	EXEC		10S	PP	FRI 5:02P	AJOB,MRS.T

2 JOBS:

0 INTRO; 0 SCHEDULED
0 WAIT; INCL 0 DEFERRED
2 EXEC; INCL 1 SESSIONS
0 SUSP

JOBFENCE= nn; JLIMIT= nn; SLIMIT= nn

When you have returned the system to its normal status and restarted suspended jobs, the backup is complete, and users can begin using the system again.

When STORE is Complete
Allowing Users Back on the System

10 Backing Up the System Configuration

This chapter describes how to use the SYSGEN utility to perform a system backup. The subjects in this chapter are:

- SYSGEN Backup Capabilities
- Running SYSGEN
- The System Backup Command
- Creating a System Load Tape
- Combined SLT and Store set

SYSGEN Backup Capabilities

A complete system backup consists of system files, user files, the file system directory, and system configuration information. `STORE` copies only system and user files and the file system directory to backup.

To back up your system configuration, use the SYSGEN utility. This creates a system load tape (SLT). A system load tape contains system configuration information, and all the system files necessary to boot and run a basic system.

SYSGEN requires that you have system supervisor (OP) capability to view system configuration data and system manager (SM) capability to save changes.

SYSGEN provides a series of command-driven user interfaces and online help facilities that describe command syntax and options. Through the interfaces, referred to as configurators, you can build new system configurations and generate a system load tape based on the new information.

SYSGEN stores configuration data in a group of files that are maintained by one or more of the five system configurators. Each configurator provides you with an interface to make changes to an independent portion of the configuration data. For example, there is an I/O configurator that allows you to change the physical makeup of the system and a Log configurator for making changes to the type and quality of system and user log files.

The base group is the group containing the set of configuration files to be read or modified by SYSGEN. The default group is the group that was specified (or defaulted to) when the system was started. If a base group is not specified, then the default group becomes the base group. The default group is generally `CONFIG.SYS`.

`CONFIG.SYS` is a permanent fileSet containing information on device classes, assigned I/O paths, assigned LDEV numbers, volume names, and various data in system files. When changes are made to the system configuration from one or more of the configurators, the files remain unchanged until you formally keep the new information through the SYSGEN `KEEP` command. You can always back out of changes until you actually keep them. You may `KEEP` the base group or to a group you specify.

Multiple configurations can be kept on disk and stored to backup through a full system backup. When you `KEEP` your changes to another group name (for example, `CONF950`), a new group is created in the `SYS` account with an identical fileset. When a system load tape is generated, however, the current configuration data is always written to the SLT as `CONFIG.SYS` regardless of what your working name on disk is for the group. This reduces confusion over which group is the current configuration or what its original name was if you suddenly have to regenerate a system. For information on changing or creating configurations groups, refer to *System Startup, Configuration, and Shutdown Reference Manual*.

Running SYSGEN

To run the SYSGEN utility, enter the following:

```
:SYSGEN basegroup,newgroup,inputfile,outputfile
```

All of the parameters are optional.

The *basegroup* parameter is the group in the account SYS that contains the configuration data you use or alter. The default is CONFIG.SYS (or the group specified at system startup). If you specify a group that does not exist, an error occurs. You can change the base group within SYSGEN using the BASEGROUP command. Refer to *System Startup, Configuration, and Shutdown Reference Manual* for additional information.

The *newgroup* parameter is the group to which you will store your data. If you do not specify this parameter, SYSGEN stores the changes to the base group. If the group you specify already exists, SYSGEN asks you whether you actually want the configuration data written to that group. To override *newgroup*, specify an alternate group with the SYSGEN KEEP command.

The *inputfile* parameter is the actual file designator of the file for command input. The formal file designator is SYSGIN. The default is \$STDIN. Use a file equation to change the command input file.

The *outputfile* is the actual file designator of the file for output requests. The formal file designator is SYSGOUT. The default is \$STDLIST. Use a file equation to redirect output.

For example, to change input to a command input file and redirect output to the line printer, enter:

```
:FILE SYSGIN=NEWIN
:FILE NEWOUT;DEV=LP

:SYSGEN,MYGROUP,*NEWIN,*NEWOUT
```

When you run SYSGEN, you receive a list of SYSGEN first-level commands:

Figure 10-1. SYSGEN First-Level Command Listing

```
SYSGEN version E.00.00 : catalog version E.00.00
FRI, JAN 7, 1994, 5:57 PM
Copyright 1987 Hewlett-Packard Co. All Rights Reserved

**warning** no NMCONFIG file in this configuration.
**warning** NMCONFIG.PUB.SYS is now the NMCONFIG file
** First level command **

      io          log (lo)          misc (mi)          spa (sp)
sysfile (sy)

      basegroup (ba)  keep (ke)          permyes (pe)      show (sh)
      backup (ta)

      clear (cl)(c)   exit (ex)(e)       help (he)(h)      oclose (oc)
      redo

sysgen>
```

NOTE The warnings refer to the network manager configuration file, which configures remote devices and terminals on the system. SYSGEN is concerned with the devices actually interfaced to the system.

The System Backup Command

The SYSGEN command that you use to make a system load tape is `TAPE (TA)`.

To create a system load tape without additional files, enter the following command:

```
sysgen>TAPE
```

or

```
sysgen>TA
```

At the end of the tape creation, you receive the following message:

```
** Boot backup generated successfully **
```

If an error occurs during the tape creation, a flashing `TAPE ERROR` message is issued to the console, describing the error.

Creating a System Load Tape

To create a system load backup that contains a new configuration, follow the sample commands as illustrated:

Figure 10-2. Creating a System Load Tape Sample

```
:SYSGEN CONFIG,NEWFIG

:SYSGEN version E.00.00 : catalog version E.00.00
FRI, JAN 7, 1994, 5:57 PM
Copyright 1987 Hewlett-Packard Co. All Rights Reserved

    **warning** no NMCONFIG file in this configuration.
    **warning** NMCONFIG.PUB.SYS is now the NMCONFIG file

    ** First level command **

    io                log (lo)      misc (mi)      spa (sp)
sysfile (sy)

    basegroup (ba)    keep (ke)      permyes (pe)  show (sh)
    backup (ta)

    clear (cl)(c)    exit (ex)(e)  help (he)(h)  oclose (oc)
    redo

sysgen>
```

Make your system configuration changes through the utility configurators. For information, refer to *System Startup, Configuration, and Shutdown Reference Manual*. Return to the first-level SYSGEN menu and keep the changes to the new configuration group:

```
sysgen>KEEP

or

sysgen>KE

Keeping to group NEWFIG
Purge old configuration (yes/no)?Y

sysgen>
```

At this point, your configuration changes have been saved to disk under the configuration group NEWFIG. If you specify NEWFIG as the base group and you issue the TAPE command, NEWFIG is written to the backup as CONFIG.SYS, overwriting the default configuration files.

Recall the new configuration and generate a system load backup:

```
sysgen>BA NEWFIG

sysgen>TA
```

If you are completely confident that you have a working configuration, exit SYSGEN.

Otherwise, recall the original CONFIG.SYS and store it under a different group name:

```
sysgen>BASEGROUP CONFIG
```

or

```
sysgen>BA CONFIG
```

```
sysgen>KEEP OLDFIG  
sysgen>
```

```
sysgen>EXIT  
END OF SUBSYSTEM  
:
```

If the system does not boot under the new CONFIG group, you can use the START GROUP=OLDFIG option under the initial system loader (ISL) to boot the system under the original CONFIG group. For more information on the ISL START command and options, refer to *System Startup, Configuration, and Shutdown Reference Manual*.

Combined SLT and Store Set

MPE/iX allows you to combine an SLT and any store filesets to make a combined SLT/Store set. This is recommended for DDS devices primarily whose large capacity allows both the SLT and store filesets to fit on a single cassette.

In a combined SLT/Store set, the files are stored after the SLT. These files can be retrieved either by a reload or a separate `RESTORE` command which will skip over the SLT.

In order to create a combined SLT/Store set backup, `SYSGEN` needs to be invoked. The option `STORE` is added to the `TAPE` command inside `SYSGEN`. The `TAPE` command looks as follows:

```
sysgen> help TAPE
tape (ta) [mode = VERBOSE|NOCHANGE|NOCONFIG|NODIAG|
           NOLOGNUM|LOGNUM]
           [dest = OFFLINE]
           [store = filesetlist string]
```

When the `STORE` option is used, files will be stored after the SLT as specified through the `filesetlist` string. The `filesetlist` string parameter is the same as the `filesetlist` parameter that is specified when using the `STORE` command. The only difference is that it specifies a set of files to be stored after an SLT as opposed to just a normal backup.

An example of how to use the `STORE` option is as follows:

```
sysgen> TA STORE = "@.pub.sys"
```

or

```
sysgen> TA STORE = " "
```

If the `STORE` option is used without specifying a `filesetlist` (2nd example), which is the same as invoking the `STORE` command without a `filesetlist`, the files that are stored will depend on the user capabilities. Since the parameter to the `STORE` option is a string, it must be enclosed by quotes.

If you want to specify other `STORE` options like `SHOW`, `PROGRESS` etc., they can be specified through the use of an indirect file:

```
sysgen> TA STORE = "indirectfile"
```

where `"indirectfile"` might contain:

```
@.@@. ;SHOW;PROGRESS=1
```

`STORE` options `TRANSPORT`, `STORESET` and `MOSET` (TurboSTORE) are invalid when a combined SLT/Store set is being created. If the store results in an error, the normal `STORE` errors will be displayed, and the user will be left in `SYSGEN`.

Storing files after the SLT will not work for a remote device. The use of the `TAPE` command to store files is not backward compatible (prior to 3.1).

When a single reel combined SLT/Store set is used to restore files, `RESTORE` skips past the SLT set automatically prior to restoring files. In a multi-reel backup, if one of the reels that contain SLT information is mounted, `RESTORE` skips past the SLT beginning with that reel, that is, reel one does not have to be mounted first in order for the `RESTORE` to work.

11 Restoring Files

This chapter describes the `RESTORE` process and options. You `RESTORE` a file by transferring it from a `STORE` backup to disk. You might need to `RESTORE` a single file, for example, when a user accidentally deletes it. You might restore all the files in an account after you have permanently archived the account offline, deleted the files, and a user needs access to the files again. After a major system failure, you might need to `RESTORE` all of your files from backups.

NOTE Restoring from sequential devices is supported in all `STORE` and `TurboSTORE` products. However, if you want to restore from parallel devices, you must have either `TurboSTORE/iX II (B5151AA)` or `TurboSTORE/iX 7x24 True-Online Backup (B5152AA)`.

This chapter describes step-by-step procedures for restoring files with `RESTORE`. The sections in this chapter are as follows:

- Preparing to Restore Files
 - Restoring Different Format Tapes
 - Referencing Files in a `RESTORE` Command
 - Issuing a `RESTORE` Command
 - Restoring from Labeled Tapes
 - Restoring from Single Devices
 - Restoring from Multiple Devices (`RESTORESET`)
 - Restoring Files from a Disk File
 - `RESTORE` Command Options
 - Monitoring `RESTORE` Command Progress
 - When `RESTORE` is Complete
-
- Do not use `RESTORE` to transfer any of the files originally on your system load tape from `LDEV 1` to a system disk other than `LDEV 1`. These files are contained in the `SYS` account and can be displayed by entering the `SY` configurator of the `SYSGEN` utility and using the `SHOW` command.
 - Do not purge or put a lockword on these files.
 - Do not use any command other than `SYSGEN` to manage the files within a configuration group. These configuration files come as a set and individual
-

files may not be moved between groups or systems. Such actions make these system files inaccessible, causing problems either immediately or at a later time.

NOTE For simplicity, the term tape is used throughout the procedures in this chapter. Substitute the appropriate media and device where applicable.

Preparing to Restore Files

Since restoring files involves the use of one or more tapes on a tape device, you must complete the following three steps before issuing the `RESTORE` command:

- Locate the necessary backup media.
- Check for duplicate filenames.
- Prepare the tape and the tape device.

Locate the Necessary File or Files

Search your tape library for the `STORE` tape(s) containing the files you want to restore. If you do not know where the tape file(s) are, you can use the `RESTORE` command to list the contents of the tape. Refer to "When `RESTORE` is Complete", later in this chapter.

Checking for Duplicate Filenames

Before you `RESTORE` a file from a `STORE` tape, check whether a disk file with the same name already exists. If a file with the same name resides in the account and group you are restoring to, the disk file is overwritten by the backup file during the `RESTORE` process by default.

CAUTION If you do not wish to lose current disk files with the same fully qualified filenames as the files you are restoring, use the `KEEP` parameter of the `RESTORE` command.

To check for duplicate filenames, use the `LISTF` command for each file you intend to `RESTORE`. For example:

```
:LISTF filename.groupname.accountname
```

Use wildcard characters to represent a set of files. Refer to "Using Wildcard Characters", in Chapter 5, "Storing Files," for details. For example, the following `LISTF` command lists all files in the `PUB` group of the `RESEARCH` account:

```
:LISTF @.PUB.RESEARCH
```

If at least one file exists, the system prints the filename at your terminal. If files matching the search criteria do not exist, the system prints the message:

```
NON-EXISTENT FILE (CIERR 907)
```

Preparing the Media and the Backup Device

Prepare the `STORE` tape following the instructions in Chapter 3, "Preparing Storage Devices." Make sure that you protect the files on the tape by removing the write ring. Mount the `STORE` tape on your tape device.

Referencing Files in a RESTORE Command

You can name a single file to RESTORE, several individual files, or all the files from one or more STORE tapes. The group of files you want to RESTORE is called a fileset. In addition to naming many kinds of filesets, you can name file subsets to exclude from the RESTORE process.

As a system supervisor (with OP capability) or a system manager (with SM capability), you can RESTORE any file in the system. To RESTORE a file that is not in your logon group and account, you must specify the complete name of the file in the form *filename.group.account*. If you do not name an account for the file, RESTORE assumes your logon account.

If you do not name a group, RESTORE assumes your current group. Thus, if you want to restore a file from your logon group and account, you need only specify its filename. RESTORE does not RESTORE a file if its group and account do not exist on your system, unless you instruct RESTORE to create accounts and groups (using the CREATE parameter). Refer to "Creating Accounts, Groups, and Creators", later in this chapter.

RESTORE uses the same guidelines and syntax to describe files to be restored as STORE. See the "Referencing Files in a STORE Command," in Chapter 5, "Storing Files."

You do not need to specify a fileset for the RESTORE command. If you omit a list of files in the RESTORE command, the default is to restore a fileset based on your capabilities, as follows:

- *@.group.account* for users
- *@@.account* for the account manager (AM)
- *@@.@* for the system manager (SM) and system supervisor (OP)

Restoring from Labeled Tapes

Restoring files from labeled tapes requires that the `RESTORE` command include a `RESTORE` file reference that was named previously in a file equation. For example:

```
:FILE TAPE7=BACKUP.JAN29;DEV=TAPE;LABEL="VOL25",ANS,02/28/90
```

To recover all files in the volume set, you would then enter the following `RESTORE` command.

```
:RESTORE *TAPE7;@.@@;restoreoptions
```

The system then requests that you mount the first volume of the volume set. For example:

```
?23:20/#S1/49/Mount tape of volumeset VOL25 (ANS)
```

After you mount the volume and it is recognized by the system, `RESTORE` continues. If the tape mounted is the first volume of the set, file recovery begins automatically. If the tape mounted is not the first volume of the set, you are prompted for the first volume. For example:

```
23:17/15/Vol XYZ (ANS) of set VOL25 mounted on LDEV# 7
```

```
23:17/#S1/50/MOUNT REEL 1 OF VOLSET VOL25 FOR RESTORE ON LDEV 7
```

NOTE Labeled tapes cannot be used to `RESTORE` with a sequential device configuration.

Restoring from Single Devices

The `FILE` command gives the tape device a filename and identifies the device. To `RESTORE` files from only one device, back reference that device in your `RESTORE` command. For example:

```
:FILE T;DEV=TAPE  
:RESTORE *T;@.OPERATOR.SYS
```

Restoring from Multiple Devices (RESTORESET)

To RESTORE files from multiple devices, use the RESTORESET parameter in the RESTORE command to reference the multiple devices. Use of the RESTORESET parameter is explained in the following pages.

Sequential Devices

You can use sequential device configuration when you have a large set of files to RESTORE from a number of tapes and you want to speed up the process. The sequential device configuration is described under "Using Tape Devices Sequentially," in Chapter 3, "Preparing Storage Devices."

To RESTORE files from sequential devices:

1. Mount the tapes on all of the devices you plan to use.
2. Use the FILE command to give each device a filename that indicates its position.
3. Use the RESTORESET parameter in a RESTORE command to reference the multiple devices.

For example, to RESTORE files from three sequential devices, use the following set of commands:

```
:FILE SEQ1;DEV=7
:FILE SEQ2;DEV=8
:FILE SEQ3;DEV=9
:RESTORE ;@.@@;RESTORESET=(*SEQ1,*SEQ2,*SEQ3)
```

NOTE Do not back reference the filename of a tape device immediately before the name(s) of the file(s) you are restoring when using the RESTORESET parameter. Instead use a placeholder for the missing parameter normally supplied when using a single tape device. The placeholder is provided by the use of the first semicolon (;) in the RESTORE command.

The above RESTORE command restores all system and user files from three sequential devices. The parentheses surrounding back-referenced filenames group the devices into a device pool.

When RESTORE has transferred all the files from the first tape on the first tape device, it immediately begins restoring files from the second tape on the second tape device. While RESTORE continues, you can unload the first tape and mount a new tape on the drive. When RESTORE finishes restoring the files on the last device, it reverts to restoring files from the first device. The process continues to cycle through the devices in the device pool until it has restored all the files you specified.

If you gave the tape devices a single filename such as:

```
:FILE T;DEV=TAPE
```

Enter the following RESTORE command to RESTORE files from a set of three sequential tape devices:

```
:RESTORE ;@.@@;RESTORESET=( *T, *T, *T)
```

You can RESTORE files that were created using parallel devices or parallel device pools in a sequential configuration. For example, to RESTORE a set of reels that have been created using parallel device pools, you must use the following logical order to mount the reels on sequential devices: A1, A2....B1, B2 and so on. (The first alphabetic character, in capital letters, indicates an original, unique parallel set; the second numeric character indicates an original sequential reel within the set.) That is, mount A1 first, followed by A2. The reel that you mount prior to B1 should be the last "A" reel.

NOTE Sequential RESTORE is not available for labeled tapes.

Parallel Devices

A parallel device or parallel device pool configuration can be used for restoring files. Therefore, you can use the same designations with the RESTORESET parameter that you can use with the STORESET parameter. Refer to Chapter 3, "Preparing Storage Devices," for information on storing to multiple devices.

If you gave the tape devices a single filename such as:

```
:FILE T;DEV=TAPE
```

Enter the following RESTORE command to copy files from a set of two parallel tape devices:

```
:RESTORE ;@.@@;RESTORESET=( *T ), ( *T)
```

Describe parallel devices in the RESTORESET parameter by enclosing each back-referenced device filename within parentheses. Use commas to separate devices.

To RESTORE files from parallel device pools, use parentheses to group the devices into pools. Separate the devices in the pool from one another with commas. Commas also separate each device pool.

For example, the following commands copy files from two parallel device pools. Each device pool contains two sequential devices. In the following example, RESTORE uses a total of four devices divided into two device pools.

```
:FILE PPA1;DEV=7  
:FILE PPA2;DEV=8  
:FILE PPB1;DEV=9  
:FILE PPB2;DEV=11  
:RESTORE ;@.OPERATOR.SYS;RESTORESET=( *PPA1, *PPA2 ), ( *PPB1, *PPB2)
```

If you gave the tape devices a single filename such as:

```
:FILE T;DEV=TAPE
```

Enter the following RESTORE command to copy files from a set of parallel device pools:

```
:RESTORE ;@.@@;RESTORESET=( *T, *T ), ( *T, *T)
```

NOTE Parallel RESTOREs can be performed only with tape sets that were created with a parallel STORE parameter. A sequential RESTORE can be used with any tape set.

Parallel RESTORE allocates every subset of the tape set (for example, STORESET=(*T1) , (*T2) creates two subsets) to distinct parallel devices or device pools, and expects a volume belonging to the subset to be mounted on its allocated device or device pool only.

When DIRECTORY option is used with RESTORESET, please restore the directories first if they are not already on the system with "RESTORE ;;DIRECTORY" command before restoring the files with the RESTORESET option

Restoring Files from a Disk File

You can restore files from a disk file created by TurboSTORE by specifying file equations for the disk files. You can restore from parallel disk files by using the RESTORESET parameter. Details on restoring files from disk files are covered in the following sections.

NOTE Restoring data from disk files created with TurboSTORE/iX 7x24 True-Online Backup is only supported on systems running MPE/iX Release 5.5 or higher. Although only the TurboSTORE/iX 7x24 True-Online Backup product can create disk files, any version of STORE or TurboSTORE running on MPE/iX 5.5 or later can restore data from disk files.

Specifying the Disk File

You direct RESTORE to restore files from a disk file created by TurboSTORE by using file equations. These file equations follow the same rules that TurboSTORE used when it created the disk files. You must specify the name of the disk file to be used in the file equation, and specify the DEV=DISC option. All other options specified in the file equation will be ignored. For example, to restore from the disk file MYBACKUP.BACKUPS.SYS, issue the following file equation:

```
:FILE DISKF=MYBACKUP.BACKUPS.SYS;DEV=DISC
```

You can then use this file equation in your RESTORE command, just as you would use a file equation for a tape device:

```
:RESTORE *DISKF;@. @. @;SHOW
```

You can change the name and location of the disk backup file from what it was when STORE created it. You then specify their new name and/or location when issuing the file equation. See the "Storing to Disk Files" section in Chapter 3, "Preparing Storage Devices," for more information on the naming convention used by TurboSTORE when it

creates disk backup files.

Restoring From Multiple Disk Files

If the backup to be restored contains more than one disk file, create only one file equation pointing to the *first* disk file. RESTORE will automatically open the subsequent disk files. Do not create any file equations for the other disk files.

Be sure that all of the disk files in a backup are present at the time of the backup. RESTORE will abort with an error if it is unable to open a disk file that is part of the backup.

All of the disk files for the backup must reside in the same group or HFS directory. RESTORE will use the fully qualified path for the file specified with the file equation to determine where the other disk files are located.

If you want to rename the disk files or change their location from where STORE originally created them, be sure that all disk files are put in the new location. Also, the *base* filename (i.e., without the ".#" extension), must all remain the same. RESTORE cannot locate disk files that are part of the same backup if they have different base filenames.

Restoring from Parallel Disk Files

If the backup was created by writing to multiple, parallel disk files with the the STORESET parameter, then files from that backup can be restored with the RESTORESET parameter. For each parallel disk file to be restored, a file equation must be issued for the first disk file in each parallel set.

For example, to restore from a two-way parallel disk backup, specify the following file equations:

```
:FILE SET1=FULL1.BACKUPS.SYS;DEV=DISC  
:FILE SET2=FULL2.BACKUPS.SYS:DEV=DISC
```

These file equations can then be used with the RESTORESET parameter:

```
:RESTORE ;@.@.SYS;SHOW;RESTORESET=( *SET1 ) , ( *SET2 )
```

As when restoring from multiple disk files, the *base* names of all disk files in each parallel set must be the same, and all files of a parallel set must reside in the same group or HFS directory. However, the names can be changed from what they were when they were created by TurboSTORE.

Unlike restoring from tape devices, when performing a parallel restore, all parallel sets in the backup must be specified in the RESTORESET command. Tape devices can use console messages to request that media from an additional parallel backup be mounted. However, when restoring from disk backups, no operator messages are issued, and if a parallel set was not specified in the RESTORE command, RESTORE will fail to restore all files on the missing parallel sets, and will abort. Always specify all parallel sets in the RESTORESET parameter.

NOTE Just as storing to serial disk files with the STORESET command is not allowed

when storing files, restoring from serial disk files using the `RESTORESET` command is not supported. You need only to specify a file equation for the first disk file in the backup, or in a parallel set. `RESTORE` will automatically find and open any additional files in the backup that it needs.

All other `RESTORE` options can be used when restoring from disk files.

Verifying Disk Backups

Disk files can also be specified for the `VSTORE` command. All of the `RESTORE` guidelines regarding file equations and the `RESTORESET` parameters apply. For example, to verify all files written to the disk file backup, `MYBACKUP`, issue the following commands:

```
:FILE DISCF=MYBACKUP  
:VSTORE *DISCF;@.@@;SHOW
```

Issuing a RESTORE Command

To RESTORE the files from the tape to the system disk, you may optionally first issue a FILE command. The optional FILE command names the tape device and describes its type. If you do not use a FILE command, RESTORE creates a default filename. You must issue a RESTORE command describing:

- The files you want to RESTORE
- The options you want to use, if any
- Back references to the filename of the tape device, unless you use the default filename.

For example:

```
:FILE T;DEV=TAPE  
:RESTORE *T;FILE1.OPERATOR.SYS
```

OR

```
:RESTORE ;FILE1.OPERATOR.SYS
```

NOTE If the back reference is omitted, the semicolon (;) serves as a placeholder and must be included.

You must use a FILE command if you want the device to be automatically allocated.

Refer to "Naming the Backup", in Chapter 3, "Preparing Storage Devices," for more information on the FILE command.

When using FOS RESTORE, you can only restore files from a single tape device, or from multiple tape devices in a serial pool. You cannot restore files from parallel devices. If you have purchased TurboSTORE/iX II (B5151AA) or TurboSTORE/iX 7x24 True-Online Backup (B5152AA), then you can use parallel devices or parallel device pools to restore files.

Restoring Different Format Tapes

RESTORE can restore files from any tape created by STORE. This includes tapes that were made using TurboSTORE only options, such as compression (COMPRESS), online (ONLINE), interleave (INTER), and parallel store sets (STORESET=). Even if you have not purchased a version of TurboSTORE, basic RESTORE can still copy the files off of the tape.

For a tape that was created by an online STORE, the original state of the file is constructed based on the written file data, plus any online log information that follows the file. The file will be restored to the state it was at the time of the attach.

You can also RESTORE files from a tape that was created with parallel store sets, even if you have only one tape device, or have fewer tape devices than were used to create the tape. Mount the media from the first store set initially. When RESTORE has finished reading files from that set, it will prompt you to mount media from the next set. It will continue to request media until all files have been restored.

RESTORE can restore files from a tape which is in MPE/VE STORE format. MPE/VE STORE format tapes are made when the TRANSPORT option is specified on an MPE/iX STORE command. When restoring from an MPE/VE STORE format tape, the STORE or TSTORE program will sense that the tape is in MPE/VE STORE format and invoke CMSTORE to RESTORE the files. However, the following options are not available in CMSTORE:

- DIRECTORY
- LISTDIR
- FCRANGE
- VOLSET
- VOL
- VOLCLASS
- COPYACD
- NOACD
- RESTORESET
- TREE
- NOTREE

RESTORE Command Options

The following sections describe the RESTORE options. The options are:

- SHOW
- ONERR
- CREATE
- CREATOR
- GID
- KEEP, NOKEEP
- OLDDATE, NEWDATE
- DIRECTORY
- LISTDIR
- PROGRESS
- FCRANGE
- FILES
- DEV
- VOL, VOLCLASS, VOLSET
- COPYACD, NOACD
- TREE, NOTREE
- PARTIALDB
- STOREDIRECTORY
- RESTORESET

For a description of RESTORESET, see the "Restoring from Multiple Devices (RESTORESET)" section earlier in this chapter.

Listing Restored Files (SHOW)

Whenever you RESTORE a set of files, the system displays the total number of files restored at your terminal. If there was an error, it also displays the names of files not restored, the reason each was not restored, and the total number of files not restored. Use the SHOW option to display the names and additional information about the files restored and to list them on your system printer as well as at your terminal.

The output listings are in the same format as STORE, which are described in Chapter 6, "STORE and TurboSTORE Command Options." Refer to that chapter for information on the possible SHOW suboptions, as well as the format and fields in the listings.

Two differences between STORE and RESTORE listing do exist, however:

- The filenames displayed in a RESTORE listing are those of the final destination filename, which may not be the same name the file has on the tape. Restore options such as GROUP, ACCOUNT, and LOCAL may affect what filename is printed.
- For the LONG listing, RESTORE has an extra field, called OLDSP. This field will contain the old spoolfile name for any spool files that are being restored. Since a spoolfile is given a new name when it is restored, this field will help you to associate new spoolfiles with their previous names.

Selecting a RESTORE Error Recovery Method (ONERROR)

As RESTORE restores files from tape, it displays the number of files restored and not restored, and the names of files not restored, if any. The RESTORE message also explains why each file was not restored. Such errors do not necessarily abort the RESTORE process. When RESTORE encounters an error, it either automatically recovers or terminates, depending upon the nature of the error and the error recovery method you choose.

The following errors always cause RESTORE to abort:

- A command syntax error.
- A disk input or output error (in the system).
- A file directory error.
- An error opening the tape file or an indirect file. Refer to "Using Indirect Files," in Chapter 5, "Storing Files.")
- An incorrectly formatted STORE tape.
- No continuation reel. You did not find a continuation reel for a multi-reel tape set.
- A device reference error. Either the specification for the device parameter is illegal, or the device is not available.

The ONERROR option of the RESTORE command lets you choose an error recovery procedure. Your options are ONERROR=QUIT, ONERROR=SKIP, or ONERROR=FULL. QUIT is the default. You do not have to specify the ONERROR option if you want the default behavior.

If you specify SKIP, RESTORE skips the file in which the error occurred and continues restoring files from the tape. If you do not specify the ONERR option, or specify ONERR=QUIT, RESTORE terminates upon encountering a tape error.

You must have Privileged Mode capabilities to specify ONERROR=FULL. This option tells RESTORE to recover as much of a file as possible if a media error is encountered when reading the file. The file is restored, if possible to the system. However, parts of the file may be missing. These *missing* parts are filled with the default fill character for the file.

For each missing piece of file data, RESTORE prints a warning message to the listing file. Also, a warning message is issued for the file, stating that it was "partially restored."

In the final count of files restored at the end of the restore operation, you get a total count of the number of files, HFS directories, and symbolic links that were partially restored. You should look carefully through your RESTORE listing when you use ONERROR=FULL, to make sure that you know which files were partially restored. These files may be corrupted, and should be inspected carefully before use. Non-flat files, such as message files or KSAM

files, are particularly likely to have experienced damage.

NOTE This option is provided only for use in the event that a media read error is preventing you from recovering an important file that can not be recovered from any other backup or data source. This option should NOT be used on a regular basis for RESTORE. It should only be used after a normal RESTORE failed to recover a file or files due to a media read error.

Restoring Files to the Correct Group, Account, and Creator (CREATE)

Files on STORE tapes belong to the same group, account, and creator that they belonged to on disk. You RESTORE files to their original group, account, or creator, or if you have the correct capabilities, you can use RESTORE command options to copy a file to your own group and account or to copy a file to a different group, account, or creator.

Creating Groups, Accounts, and Creators

If a file's account, group, or creator has been deleted from your system after storing the file to tape, you can recreate it as you RESTORE the file from tape using the CREATE option. RESTORE sets account, group, and user capabilities to their default values when it creates them. You must have the necessary capabilities to create a new group or account. That is, you must have system manager (SM) or system supervisor (OP) capability to create a new account. You must have system manager (SM), system supervisor (OP), or account manager (AM) capability to create a new group.

For example, you stored all files in the account FEBRECS to tape on the first of March and then purged the account, its users, and its groups from the system. Several months later, a user asks you to RESTORE the files in the FEBRECS account. Using the CREATE option, you can recreate the account, groups, and creators as you RESTORE the files. For example:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@.FEBRECS;CREATE=GROUP,ACCOUNT,CREATOR
```

Or alternatively, you could enter:

```
:RESTORE;;CREATE
```

which will by default RESTORE all files and create all missing structures.

CAUTION Using this method, the CREATE option creates groups, accounts, and users with default capabilities and access rights (those that exist for the group, account, and user to which you are restoring, not those that exist on the tape).

Restoring Files to Your Group and Account

Regardless of the group, account, or creator from which files were stored, you can RESTORE files into your own group and account if you have read access to the files on the tape, or system manager (SM), system supervisor (OP), or account manager (AM) capability.

"Read access" implies that if the files were restored to the groups and accounts from which they came, and those groups and accounts had default access capabilities, you would be able to read the files on disk, such as with FCOPY.

Use the LOCAL option. For example, enter:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@@;LOCAL
```

The files are restored to your logon group and account with your logon user name as the creator.

Restoring a File to a Different Group, Account, or Creator

Similarly, use the GROUP and ACCOUNT options to switch a file's group and account as you RESTORE it from backup. You can use either option alone, or use both together, but you cannot use either GROUP or ACCOUNT with LOCAL. You must have system manager (SM) or system supervisor (OP) capability to switch a file's account. You must have system manager (SM), system supervisor (OP), or account manager (AM) capability to switch a file's group.

NOTE A user without the capabilities described above can RESTORE a file to a different account or group if the following conditions are met:

- The group, account, and file level access are available to the user.
- The file has no lockword, or if it does, the user knows the lockword.
- The file is not privileged.

The commands displayed below RESTORE the files that belong to the PUB group of the SMITH account on tape to the PUB group of the JONES account on disk.

```
:FILE T;DEV=TAPE
:RESTORE *T;@.PUB.SMITH;GROUP=PUB;ACCOUNT=JONES
```

The CREATOR option lets you change a file's creator as you RESTORE the file from tape. For example, the following command restores the files in the PUB group of the SMITH account on tape to the PUB group of the JONES account on the system disk, changing the creator to MARTY:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.PUB.SMITH;GROUP=PUB;ACCOUNT=JONES;CREATOR=MARTY
```

You must name a user that exists in the account as the new creator, or specify CREATE to cause the user to be created in the directory. If RESTORE does not find the user name, it does not RESTORE the file. If you use the CREATOR option without specifying a user name, RESTORE restores the file only if the tape file's creator exists in the file system directory.

NOTE The `CREATOR` option cannot be used when the `LOCAL` option is used.

Changing a File's GID (GID)

If you need to change a file's group ID, or GID, you may specify the `GID` option. The `GID` option takes an optional file group name, for example:

```
;GID=MANAGER
```

If the file group name is specified, all files being restored will have their GID changed to the specified GID. If the file group name is not specified, the GID present on the backup will be preserved. This overrides any change in GID that may occur due to the `LOCAL` or `ACCOUNT` options.

Overwriting or Retaining Disk Files (KEEP, NOKEEP)

By default, `RESTORE` replaces disk files with the same fully qualified filename as the file you are restoring from backup. Therefore, especially if you are restoring a large number of files, to be sure that you do not overwrite an important file or files, use the `KEEP` parameter of `RESTORE`.

For example, your `STORE` tape might contain several files, some of which have names that are the same as files on disk. To `RESTORE` only those backup files with names that do not duplicate disk file names, use the following commands:

```
:FILE T;DEV=TAPE  
:RESTORE *T;@.@@;KEEP
```

Using `KEEP` in the command above tells the system not to replace the files already on disk with files that have identical names on tape. You successfully `RESTORE` to disk all files on the backup that do not have the same names as files already on disk. You do not overwrite existing files on the disk.

To explicitly require `RESTORE` to overwrite disk files with the same fully qualified filenames, use the `NOKEEP` option in your `RESTORE` command. For example:

```
:FILE T;DEV=TAPE  
:RESTORE *T;@.@@;NOKEEP
```

Specifying Modification and Last Access Date (OLDDATE, NEWDATE)

When you `RESTORE` a file or files, you can choose either to retain the creation, modification, access, and state change dates and times stored with the file on the backup, or you can choose to change the file's dates and times to the date you restored the file.

To retain the creation, modification, access, and state change dates and times in the file label on the backup, use the `OLDDATE` option of `RESTORE`. For example:

```
:FILE T;DEV=TAPE  
:RESTORE *T;@.@@;OLDDATE
```

To change the dates and times to the date you restored the files, use the `NEWDATE` option. For example:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@@;NEWDATE
```

You will probably want to use `NEWDATE` when you `RESTORE` archived files, so that you do not quickly archive them again.

Restoring Directory Information (DIRECTORY)

Use the `DIRECTORY` parameter to restore directory information from backup. To do this, you must have system supervisor (OP) or system manager (SM) capability. All system and volume set directories located are restored. For example, to `RESTORE` all files and all directories from a backup, enter:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@@;DIRECTORY
```

If you use the `LOCAL`, `ACCOUNT`, `GROUP`, `CREATOR`, or `CREATE` parameter with `DIRECTORY`, `RESTORE` does not create or change accounts, groups, or users for files that belong to directories on tape. The `LOCAL`, `ACCOUNT`, `GROUP`, `CREATOR`, and `CREATE` parameters will create accounting structures only for the files that do not belong to the directories you restore.

All HFS directories, and directories from any other volume sets that were stored, are also restored when `DIRECTORY` is specified.

NOTE When `RESTORESET` option is used with `DIRECTORY`, please restore the directories first if they are not already on the system with "`RESTORE ;;DIRECTORY`" command before restoring the files with the `RESTORESET` option.

Finding Out What Is on Your Tape (LISTDIR)

Use the `LISTDIR` parameter of the `RESTORE` command to display information from the tape directory and tape label without restoring any files. The tape creation type, record size, and any files that match your fileset list display. `LISTDIR` may not be specified with any other parameter except `DIRECTORY`, `TREE`, and `NOTREE`.

NOTE The `LISTDIR` parameter works only using native mode `STORE` tapes and not with tapes created for MPE V/E using the automatic `TRANSPORT` mode of the `STORE` command.

The following example shows a sample `RESTORE` command and output display format using the `LISTDIR` parameter:

Figure 11-1. Sample RESTORE Command and Output

```

:RESTORE *T;@.PUB.SYS;LISTDIR;SHOW

          TURBO-STORE/RESTORE  VERSION  C.55.05  B5152AA
          (C) 1986 HEWLETT-PACKARD CO.
WED, JUN 26, 1996,  5:07 PM

MPE/iX MEDIA DIRECTORY

MEDIA NAME      : STORE/RESTORE-HP/3000.MPEXL
MEDIA VERSION   : MPE/iX 08.50 FIXED ASCII
MEDIA NUMBER    : 1

MEDIA CREATION DATE
WED, JUN 26, 1996,  4:46 PM

STORE  w@.pub.sys;*t;show

MEDIA CREATED WITH THE FOLLOWING OPTIONS

MEDIA RECORD SIZE : 32768
INTERLEAVE DEPTH  : 1

FILENAME GROUP  ACCOUNT  CREATOR          MEDIA   SET
WAIT   .PUB     .SYS      MANAGER .SYS    >=     1
WC     .PUB     .SYS      MANAGER .SYS    >=     1
WELCOME .PUB     .SYS      MANAGER .SYS    >=     1
WHAT   .PUB     .SYS      MANAGER .SYS    >=     1
WILL   .PUB     .SYS      MANAGER .SYS    >=     1
WMLOG  .PUB     .SYS      MANAGER .SYS    >=     1
WORDTEXC.PUB   .SYS      MANAGER .SYS    >=     1
WORDTEXT.PUB   .SYS      MANAGER .SYS    >=     1

```

If any HFS syntax files exist on the media being examined, the resulting listing will show the filenames in HFS syntax. The names will be displayed at the end of the line, similar to the HFS format for the `SHOW` option.

Displaying Progress Messages (PROGRESS)

Use the `PROGRESS` parameter to display `RESTORE` command progress messages at regular intervals. For example, to display progress messages every five minutes, use the following command:

```

:FILE T;DEV=TAPE
:RESTORE *T;@. @. @;*T;PROGRESS=5

```

If you use the `PROGRESS` option alone, without specifying an interval, `RESTORE` displays status messages every minute.

If `RESTORE` is being run from a MPE session, the progress messages will be sent to the standard list (`$STDLIST`). Otherwise, the messages will go to the system console.

Restoring Files with Certain File Codes (FCRANGE)

Use the FCRANGE parameter to restore only files with certain file codes. MPE/iX file codes distinguish different types of files. You can select up to eight file code ranges to restore. For example, files with codes 1100, 1101, and 1102 are HPWORD files. Files with codes 1152 and 1153 are SLATE files. To restore all HPWORD and SLATE files from a backup, enter:

```
:FILE T;DEV=TAPE  
:RESTORE *T;@. @.;FCRANGE=1100/1102,1152/1153
```

NOTE The FCRANGE parameter is not valid when restoring from TRANSPORT tapes.

Specifying a Maximum Number of Files to Restore (FILES)

By default, MPE/iX restores a maximum of 4000 files at a time. If you are restoring more than 4000 files using the automatic TRANSPORT mode, use the FILES= parameter to specify the maximum number of files to be restored. If you do not use the TRANSPORT mode, the FILES= parameter is not necessary. If the FILES= parameter is present with the TRANSPORT mode, it is ignored. For example, to set the maximum number of files restored to 6000 from a TRANSPORT tape, you might enter the following:

```
:FILE T;DEV=TAPE  
:RESTORE *T;@. OPERATOR. SYS;FILES=6000
```

Restoring Files to a Specific Device (DEV)

Use the DEV parameter to specify the device where you want to RESTORE files. Use an LDEV number or a device class to indicate the device. For example, the following command restores the file FILE1 to the disk with logical device number 2.

```
:FILE T;DEV=TAPE  
:RESTORE *T;FILE1. OPERATOR. SYS;DEV=2
```

If you name a device class, RESTORE allocates the files to any of the home volume set's volumes within that class. If you name a specific logical device, RESTORE restores the file to that device only if the device is a system disk.

If you do not specify a device with the DEV parameter, RESTORE tries to RESTORE a file to the logical device compatible with the device type and subtype specified in the file's label and the device type and subtype of the mounted home volume set. If it cannot find such a device, RESTORE tries to RESTORE the file to a device with the device class specified in the file's label and the home volume set. If it cannot find a device with the appropriate device class, RESTORE tries to RESTORE the file to any member of the home volume set; if it cannot, it does not RESTORE the file.

You cannot use DEV with VOLSET, VOLCLASS, or VOL. See the following section.

Restoring Files to Specific Volumes (VOL, VOLCLASS, VOLSET)

Use the VOL, VOLCLASS, and VOLSET parameters to RESTORE files to a particular volume, volume class, or volume set.

Use the `VOLSET` parameter to reference a particular volume set. If there is no room in the volume set, `RESTORE` does not `RESTORE` the file. For example:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@@;VOLSET=PRIVATE_VOL_A
```

Use `VOLCLASS` to reference a particular volume class. If there is no room in the volume class you indicate, `RESTORE` restores the files to a volume within the volume set. If there is no room in the volume set, `RESTORE` does not `RESTORE` the files. For example, to `RESTORE` files to the volume class `CLASS_B` within the `PRIVATE_VOL_A` volume set, enter:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@@;VOLSET=PRIVATE_VOL_A;VOLCLASS=CLASS_B
```

Use `VOL` to reference a particular volume. If you do not indicate a volume class or volume set, the volume you indicate must be part of the system volume set. If there is no room on the volume you name, `RESTORE` restores the file to another volume within the volume class. If there is no room within the volume class, `RESTORE` restores the file to a volume within the volume set. As an example, to restore files to the volume named `VOL_C` within the `PRIVATE_VOL_A` volume set, enter:

```
:FILE T;DEV=TAPE
:RESTORE *T;@.@@;VOLSET=PRIVATE_VOL_A;VOL=VOL_C
```

You cannot use `VOL`, `VOLCLASS`, or `VOLSET` together with the `DEV` parameter (see previous section). If a file's group and account do not exist on the volume specified with `VOLSET`, the file is restored to the volume set where its group and account do exist. Refer to *Volume Management Reference Manual* (32650-90045) for more information on using volumes, volume classes, and volume sets.

NOTE The `VOL`, `VOLCLASS`, and `VOLSET` parameters are not valid for restoring from `TRANSPORT` tapes.

Modifying ACDs (COPYACD, NOACD)

When restoring files, the `NOACD` option can be used to prevent the ACDs on the backup from being restored. This will cause the files being restored to have less security than when they previously existed on the system. Using `NOACD` can be useful when the ACDs that exist on the backup refer to users or accounts that do not exist on the current system. Specifying `COPYACD` (which is the default), will cause all ACDs associated with files on the backup to be restored to the system.

Specifying HFS Files (TREE, NOTREE)

By default, `RESTORE` uses the last character of an HFS filename to determine if the specified file or fileset should be scanned recursively to include all files below the specified wildcards or directories. If a filename ends in a slash (/), then it will be scanned recursively to include all files below it in the hierarchical directory. Otherwise, the file or fileset is not scanned recursively, resulting in a horizontal cut at its level in the hierarchical directory. Using `TREE` and `NOTREE` can override that default behavior.

If `TREE` is specified, ALL files and filesets are scanned recursively, regardless of their ending character. If `NOTREE` is specified, then NO files or filesets are recursively scanned,

resulting in a horizontal cut in the directory structure.

Restoring Part of a Database (PARTIALDB)

By default, RESTORE does not allow you to restore only part of a TurboIMAGE or ALLBASE/SQL database. The root file or DBE file and all dataset files must be specified, either by specifying just the root or DBE file, or by specifying a wildcard that includes all files in a database. However, if you need to restore only certain files from a database, then you can specify the `PARTIALDB` option to override the default behavior.

WARNING **Database corruption may result if not all database files are stored and restored. Be sure that you only want to store certain database files before overriding the default behavior with the `PARTIALDB` option.**

Refer to the "Storing Database Files Using `PARTIALDB` or `FULLDB`" section in Chapter 6 for more information on using the `PARTIALDB` option.

NOTE Quiescing an ALLBASE/SQL database and storing the ALLBASE/SQL database by TurboSTORE/iX 7x24 True-Online Backup when the DBEnvironment is in normal use, has the same effect as using the `STOREONLINE` command of SQLUtil. For a 7x24 true-online backup of ALLBASE/SQL to be of use for rollforward recovery, use the same procedures you would use when doing the `STOREONLINE` commands. (Refer to Chapter 7, "TurboSTORE/iX 7x24 True-Online Backup," for more information on quiescing.)

Restoring using a disk directory file (STOREDIRECTORY)

The `STOREDIRECTORY` option can be used to specify which store disk directory file to be used when restoring files from a backup. The backup must have been created using the `STOREDIRECTORY` option. See Chapter 6 for more information on creating backups with the `STOREDIRECTORY` option.

If this option is not specified, then RESTORE will look in the default location for a disk directory file for this backup. The default file that it will look for is:

```
/SYS/HPSTORE/store_dirs/store_yymmdd_hhmmssstt_pin##_day
```

The date, time, and pin number used to create this name are read from the store label of the media mounted for a backup. If the disk directory file exists on the system, RESTORE will read the media directory information from that file instead of from the backup media. RESTORE will then skip the media directory on the backup media and begin to restore files.

When restoring backups created with a 7x24 true-online backup sync point at the end of the backup, it is particularly important that a disk directory file exist on the system where the files are being restored. When restoring sync-at-end backups, RESTORE needs to know which files need to have after image log data applied to them, before it restores any

files. This information is stored in the disk directory file and in a media directory located near the end of the backup media.

It is much faster and more convenient for you to read the after image log information from the disk directory file. Otherwise, RESTORE needs to have you mount the backup media that contains the "final" media directory before it can restore any files.

If a disk directory file for the backup exists, but it is not located in the default location, the STOREDIRECOTRY (or STOREDIR) option can be specified to tell RESTORE where the disk directory file is located. The filename must be specified, and can be in MPE or HFS format. If it is not fully qualified, it will be qualified with the CWD.

For example, to tell RESTORE to use the file MYDIR as the disk directory file, issue the command:

```
:RESTORE *T;@.@@;SHOW;STOREDIR=MYDIR
```

or

```
:RESTORE *T;@.@@;SHOW;STOREDIR=/SYS/HPSTORE/MYDIR
```

If RESTORE cannot open the specified disk directory file, or it cannot open the disk directory file in the default location, it will use the STORE directory from the backup media. No warning message will be displayed if RESTORE cannot locate a disk directory file.

If the STOREDIRECOTRY option is used with the LISTDIR option, RESTORE can display a list of all files in the specified backup without requiring that the backup media be mounted or present on the system. This allows you to determine what files are in a backup, even if the backup media is not available. If the STORE disk directory file is located in the default location, then the full pathname of that file needs to be specified with the STOREDIRECOTRY option in order to get a listing of all files in that backup.

For example, to see what files were in the 5/12/95 backup:

```
:RESTORE *T;@.@@;LISTDIR;&  
STOREDIR=/SYS/HPSTORE/STORE_DIRS/store_19950512_22002025_pin20_fri
```

If you keep your disk directory files in a different location, then the alternate name can be specified. A filename which is a symbolic link to the default location can also be specified:

```
:RESTORE *T;@.@@;LISTDIR;STOREDIR=LASTFULL.HPSTORE.SYS
```

Monitoring RESTORE Command Progress

After entering a RESTORE command, you see the following message at the console telling you that the RESTORE process has begun:

```
TURBO-STORE/RESTORE  VERSION  C.55.05  B5152AA  
(C) 1986 HEWLETT-PACKARD CO.
```

```
WED, JUN 26, 1996,  5:07 PM
```

Responding to Tape Requests

If your system is not configured to automatically assign devices, you see a tape request. For example:

```
?9:08/#S25/43/LDEV# FOR "T" ON TAPE (NUM)?
```

If you have previously issued a FILE command that included the tape device's LDEV number, the tape request asks you to confirm that number. For example:

```
?14:57/#S25/43/IS "T" ON LDEV #8 (Y/N)?
```

See Chapter 7, "Handling Media Requests," for information on how to respond to tape requests.

Replying to Tape Requests for a Multiple-Device RESTORE

If you are restoring files from several devices, the system sends you a tape request for each device. The order that you receive tape requests corresponds to the order that you listed the devices in your RESTORE command.

If you gave your tape devices unique filenames, you receive a tape request for each unique filename. See Figure 3-1., in Chapter 3, "Preparing Storage Devices," to determine which LDEV to assign to each name.

If you gave your tape devices a common filename, you receive several identical tape requests referring to the common filename, one for each device. For example, the first tape request in a sequential RESTORE asks for the LDEV number of the first device. The second request asks for the LDEV number of the second device, and so on.

Mounting Additional Tapes

If you are restoring a large set of files, the set may extend on to several tapes. If you have more than one tape to mount on a tape device, RESTORE immediately prompts you to mount the next tape after it finishes restoring files from the previous tape. You see messages similar to the ones below:

```
15:28/#S415/59/MOUNT NEXT MEDIA FOR STORE ON LDEV 8  
15:29/#S415/59/LDEV #8 NOT READY
```

To mount another tape, perform the following steps:

1. Unload the first tape following the instructions in Chapter 3, "Preparing Storage

Devices."

2. Locate the second tape, prepare it, and mount it. You do not have to issue another `RESTORE` command to respond to a tape request. The system automatically continues restoring files. If you leave the console while restoring files, other messages can cause mount requests to scroll off of the console screen.
3. Continue monitoring the `RESTORE` process. Mount new tapes as the system requests them.

Using the PROGRESS Option

If you specified the `PROGRESS` option of `RESTORE`, you can monitor the progress of your recovery by looking at the progress messages issued by `RESTORE`. If `RESTORE` is being run from a session, then the progress messages will be issued on `$STDLIST`. If `RESTORE` is being run from a job, then the progress messages will be displayed on the system console.

When RESTORE is Complete

If you specify the SHOW parameter, as RESTORE proceeds, it lists each file restored at your terminal. For example:

Figure 11-2. Listing Restored Files

```
:RESTORE *T;FILE1.OPERATOR.SYS;SHOW  
  
WILL RESTORE          1 FILES;    NUMBER OF FILES ON TAPE =  1  
  
  FILENAME GROUP    ACCOUNT  VOLUME R'TIONS  SECTORS CODE    REEL  
  
  FILE1   .OPERATOR .SYS      DISC           :C      4      1  
  
FILES RESTORED: 1
```

When you see that all the files you named are restored, remove all tapes from your tape devices. Check that the label on the tape properly describes it, and relabel it if necessary.

If you used the SHOW=OFFLINE option, the system prints the list of files restored on the system printer. Retrieve the report from the printer and keep it for your records.

Return your tapes to your library.

12 Using Backup Media to Recover from a System Failure

This chapter describes how to recover from a major system failure. The sections in this chapter are:

- Reloading Your System Configuration
- Restoring Files from Partial Backup Tapes
- Restoring Files from Full Backup Tapes

In the event of a major system failure, you may need to reload your system from backup media. Regardless of which method you used to back up your system, your backups contain three important types of information:

- Your operating system (native library, system library) configuration.
- The file system directory.
- User and system files.

Depending upon the nature of the system failure, you might need to reload any one of these components or all three. You reload your system configuration from your system load backup, and you reload your file system directory and system and user files from your `STORE` backups. Always consult your Hewlett-Packard representative before you completely reload system information from backup media. To completely reload your system from backup media, perform the following steps:

1. Reload the system configuration from your backup system load tape. Refer to the instructions below.
2. Use the `RESTORE` command to copy system and user files and the file system directory from your most recent backups to disk. Follow the instructions below.
3. If your most recent backups are from a partial backup, use `RESTORE` to copy files from your most recent full backups to disk, keeping the files you have already restored from the partial backup. Follow the instructions below.

Reloading Your System Configuration

If your most recent system configuration files are from backups, do the following to retrieve your system configuration:

1. Install your latest system load backup.
2. Start the system. Only the system master volume on LDEV 1 is initialized. Other volumes must be initialized using `VOLUTIL`. This can be done with the following set of commands:

```
:VOLUTIL  
volutil: NEWVOL volumesetname:membername LDEV=nn
```

The *volumesetname* for the system volume set is `MPEXL_SYSTEM_VOLUME_SET`. Issue the `NEWVOL` command for each member of the volume set. None of the drives can be written to until they have been initialized. Refer to *System Startup, Configuration, and Shutdown Reference Manual* for more detailed instructions.

3. Make any changes to the `SYSGEN` configuration and reboot the system. (This step is optional.)
4. Restore your partial and full backups using the `KEEP` and `DIRECTORY` parameters of the `RESTORE` command. These procedures are described in the following sections.

Refer to *System Startup, Configuration, and Shutdown Reference Manual* for additional information.

Restoring Files from Partial Backup Tapes

Use the `RESTORE` command to copy system files, user files, and the file system directory from your `STORE` backups to disk.

To restore system and user files from your most recent backups, perform the following steps:

1. First gather all the units of backup media from your most recent backup.
2. Mount the backup(s) on the backup drive(s) and bring the drive(s) online.
3. After a system failure, only LDEV 1 is available as the system volume set. You must use the `VOLUTIL` program in order to initialize your other volume sets, if you have not done so already. This can be done with the following set of commands:

```
:VOLUTIL  
volutil: NEWVOL volumesetname:membername LDEV=nn
```

The *volumesetname* for the system volume set is `MPEXL_SYSTEM_VOLUME_SET`. Issue the `NEWVOL` command for each member of the volume set. For more detailed instructions, see the *Volume Management Reference Manual*.

4. Issue the following `FILE` and `RESTORE` commands to copy the files from the backup media.

```
:FILE T;DEV=TAPE  
:RESTORE *T;@.@@;KEEP;OLDDATE;DIRECTORY
```

5. Monitor the console for tape requests.
6. Mount additional backup media when necessary until you have copied all files from all backups.

Restoring Files from Full Backup Tapes

If your most recent backups are from a partial backup, you need to reload your remaining files from your most recent full backup.

To restore system and user files from your most recent full backups:

1. Gather all the backup media from your most recent full backup.
2. Mount the backup(s) on the backup devices and bring the devices online.
3. Issue the following `FILE` and `RESTORE` commands to copy the files from the backup device. The `SHOW=OFFLINE` option, if used, lists the files you restored on the system printer.

```
:FILE T;DEV=TAPE  
:RESTORE *T;@.@.@;SHOW=OFFLINE;KEEP;OLDDATE
```

4. Monitor the console for tape requests.
5. Mount additional backup media when necessary until you have copied all files from all backups.

13 Disaster Recovery

This chapter discusses methods and planning for disaster recovery. The sections in this chapter are:

- Recovery Facilities
- Planning for Disaster Recovery
- Testing Your Recovery System

In this age of highly computerized business, some of your business information may be available only through your computer systems. Very often the information we rely on is only available through our computers. Regularly backing up the information stored on your computer system protects that information from a system failure. Your backups and your computer system, however, are not indestructible and can be damaged or destroyed by smoke, noxious fumes, explosion, flood, earthquake, sabotage, and many other natural and man-made problems. Without access to your computer system, you may lose track of sales, production, existing customers, and accounting information.

This chapter discusses strategies for protecting computer operations from a variety of disasters. Planning for a disaster and identifying the best means for recovering with a minimum of delay is the key to disaster recovery. Costs of disaster recovery systems vary greatly depending upon the alternative you select and the features you require. However, if a properly designed disaster recovery system saves your business, you will want it at any cost.

Recovery Facilities

The best way to prepare for an unpredictable disaster is to have a recovery facility available. A recovery facility contains a backup system onto which you can reload your files from backups. You can use the backup system for all of your information processing needs until your system is repaired or replaced. There are several types of recovery facilities. You should select the type that best meets the needs of your data center.

Private Recovery Facilities: Cold Sites and Hot Sites

A private recovery facility is owned and maintained by your company. It is usually located at a facility close to, but separate from, your data center. There are two types of private backup sites, cold sites and hot sites.

Cold Sites

A cold site is a facility with sufficient electrical power, air conditioning, and telecommunications equipment to support a computer system, but it does not contain an actual computer system. If a disaster occurs, you must obtain the computer and the necessary peripherals to support your information processing requirements. The cold backup site offers guaranteed access at a relatively low cost, but requires considerable time to obtain, install, and test a new computer system.

Hot Sites

A hot site is a cold site with an operational computer system similar, if not identical, to the one you intend to protect. If a disaster occurs, you can reload your information from the backup system at the hot site. This alternative is the quickest, but most costly, method for recovering from a disaster. Because of the cost of maintaining a complete backup system, companies with hot sites sometimes begin using the backup system for regular information processing.

Mutual Backup Arrangements

Another recovery site option is a mutual backup arrangement. A mutual backup arrangement is an agreement between two companies with similar computer systems to provide the use of their Data Centers to one another should a disaster occur. Usually the two businesses are located close to one another. Some companies seek partners in different industries to eliminate competition or protect confidentiality. A mutual backup arrangement can provide both a quick and cost-effective recovery facility, but may not work well if either party is unwilling to disrupt its information processing to accommodate the other party. Contractual and legal problems sometimes arise in this type of agreement.

Hewlett-Packard Disaster Recovery Services

Hewlett-Packard offers a full range of disaster recovery services. Contact your local Hewlett-Packard representative for more information.

Commercial Backup Companies

Commercial backup companies offer hot and cold backup sites identical to private hot and cold sites except that the site is owned and operated by a company in the disaster recovery business. A commercial backup site is less costly than a private backup site, but has the disadvantage that the clients of the backup company could potentially compete for backup resources at the same time. For example, in the event of a widespread natural disaster such as an earthquake or a flood, many businesses may need the backup facility at the same time.

Computer Service Bureaus

Some companies use a Computer Service Bureau in the event of a disaster. Although such an arrangement is generally available, it is often costly, restricts computer use, causes capacity and compatibility problems, and ignores special computer requirements.

Planning for Disaster Recovery

Selecting a recovery facility is only part of a comprehensive disaster recovery program. You should prepare and implement a comprehensive disaster recovery plan. Your plan should detail resource requirements and procedures for each step in the recovery process.

The Recovery Facility

Locate your recovery facility close to the computer system you want to protect, but not in the same building or group of buildings. It should be in a different power distribution system. If you think your system may be damaged by a major earthquake, locate the recovery facility in another region of the country. The facility should have adequate security, provided by either security guards or controlled system access.

System Compatibility

The disaster recovery computer system must be compatible with your Data Center computer system. The main memory, disk drives, backup drives, and printers should allow you to perform necessary recovery operations without requiring you to reconfigure your system. The recovery computer system should also be expandable to meet your future needs.

The number of telephone lines into the disaster recovery facility should be sufficient to support recovery operations. Telephone lines should be either of voice or data quality, depending on your requirements. Data communications equipment must be compatible with your data transmission requirements. Pay attention to required baud rates and other special requirements.

Name representatives from your Data Center and the disaster recovery facility who will keep in frequent contact to ensure continued compatibility between computer systems. They ensure that changes made in either system do not affect the disaster recovery capability, and that upgrades in the Data Center computer system are quickly reflected in the disaster recovery computer system.

Disaster Notification and Response Time

A disaster can occur any time without warning. Your recovery center should be available 24 hours a day. The time from when you notify the recovery center to the time that your backup system is up and running is called response time. Response time requirements vary depending upon the business situation. A response time of four hours is normally considered a quick turnaround.

Only authorized personnel should notify the disaster recovery facility to initiate disaster recovery. They should identify themselves to the disaster recovery facility by a prearranged codeword; and, of course, the disaster recovery facility personnel must act quickly and report promptly for recovery operations.

Storage of Backup Tapes

If you keep your regular backups at your Data Center, and the Data Center is destroyed by fire, the backups are also destroyed. Create a duplicate set of backups to keep at a suitable off site storage facility. In this way, you have a set easily accessible for routine problems and a set available in case of a major disaster. Arrange to have around-the-clock access to the backups, and to have them quickly delivered from the disaster recovery facility when necessary.

Testing Your Recovery System

When you initiate a disaster recovery service contract, use backups to test for system compatibility problems and bugs in your disaster recovery system. As time passes, perform periodic tests (for example, every quarter) to ensure that your systems remain compatible and your disaster recovery procedures work well. Also test your disaster recovery system whenever you change it. You can either notify the facility in advance, to minimize costs and interruption of the recovery facility, or you can do a full-scale simulation. Either way, prepare a test plan and have it approved by all parties.

A Error Messages

This appendix describes how to interpret and apply `STORE`, `RESTORE`, and `VSTORE` error message information. The sections in this appendix are:

- Determining the Cause of `STORE`, `RESTORE` and `VSTORE` Errors
- Using the `STORE_ERROR` Variable

Determining the Cause of STORE, RESTORE and VSTORE Errors

Some STORE, RESTORE and VSTORE errors cause the program not to store, restore, or verify a particular file. Others cause the program to stop running. When STORE, RESTORE or VSTORE cannot transfer files that you named in your command, it lists the names of those files on your terminal as part of its standard progress listing.

When STORE or VSTORE aborts because of an error, you see specific error messages followed by a general error message (which is the same for both commands):

```
STORE ABORTED BECAUSE OF ERROR
```

When RESTORE aborts because of an error, you see specific error messages followed by a general error message:

```
RESTORE ABORTED BECAUSE OF ERROR (CIERR 1091)
```

You can use the STOREJCW job control word (JCW) to determine the cause of an error and a proper recovery procedure. Each time you issue a STORE or RESTORE command, the system updates the value assigned to STOREJCW.

When STOREJCW is equal to 0, the command executed successfully. STOREJCW values greater than 0 alert you to the cause of an error.

NOTE Issuing the VSTORE command also results in the system updating the value assigned to STOREJCW.

Viewing the Value of STOREJCW

To view the value assigned to STOREJCW, use the following SHOWVAR command:

```
:SHOWVAR STOREJCW
```

The system reports the current value. For example:

```
STOREJCW=4
```

Determining the Cause of an Error

The value of STOREJCW after you issue a STORE, RESTORE, or VSTORE command tells you the cause of any errors in storing, restoring, or verifying the files you named. To determine the cause of an error:

1. Use the SHOWVAR command to view the value of STOREJCW.
2. Look up that value in Table A-1.. In addition to the cause of an error, <Undefined

Cross-Reference> displays recommended recovery procedures.

Table A-1. STOREJCW Values

Value	Cause	Recovery Procedure
0	No errors.	None.
1	Syntax error.	Correct and reenter command.
2	Error in opening system resource allocation.	Purge some files or perform disk space recovery.
3	Error in opening an indirect file.	Does the indirect file you named exist? Do you have adequate access to it? Is it in use? Is the redirect file an EDIT file?
4	Error in opening the file.	Is the FILE command describing the correct device? Is the media on the device you named? Unload and reload the media. If you are restoring or verifying files, check whether or not you have adequate access to them.
5	Error scanning the files.	Make sure the file names are syntactically correct and correctly spelled.
6	Error in actually storing, restoring, or verifying files.	Enter the command again. After a second failure, contact your system manager.
7	A file you attempted to store does not exist on disk, or a file you attempted to restore or verify does not exist in the STORE files, or the system could not find files that matched the pattern you specified (for example, FEB@).	Did you name the files correctly in the STORE, RESTORE, or VSTORE command? Do the files you want to store, restore, or verify exist?
8	A file that you tried to store, restore, or verify was in use, or unable to quiesce a certain TurboIMAGE database. See the SSTDLIST for further information.	Wait until the file is no longer in use. Enter the command again. STORE may have timed out trying to quiesce an open transaction.
9	Media Manager Service error encountered.	Check the STORE listing for the media manager error message. Look up the message in Appendix A of the <i>Magneto-Optical Media Manager User's Guide</i> and follow the suggested action.

A complete listing of all possible STORE and RESTORE errors, along with the probable cause and recommended action can be found in the *MPE/iX Error Messages Manual Volumes I, II, and III*.

Using the STORE_ERROR Variable

When STORE, RESTORE, or VSTORE aborts because of an error, the Command Interpreter (CI) variable STORE_ERROR is set to a non-zero value. This value can be used to look up more detailed cause and action error text in the *MPE/iX Error Messages Manual Volumes I, II and III*. If STORE exits successfully, then STORE_ERROR is set to 0.

You can look up STORE errors in the CI error catalog by entering:

```
:HELP SR####
```

where #### is the error message from STORE. For example,

```
:HELP SR1003
```

```
MESSAGE FILE NAME LENGTH TOO LONG.  A NAME PART CANNOT EXCEED  
8 CHARACTERS IN LENGTH (S/R 1003)
```

```
CAUSE      Same as message content.
```

```
ACTION     Retype the command and include a valid file name.
```

B Command Syntax

This appendix lists the `STORE`, `RESTORE`, and `VSTORE` command syntax. The parameters, options, and variables are described in Chapter 6, "STORE and TurboSTORE Options" and Chapter 11, "Restoring Files." The sections in this appendix are:

- `STORE` Syntax
- `RESTORE` Syntax
- `VSTORE` Syntax

STORE Syntax

```
STORE [filesetlist][;storefile][;option[;...]]
```

where *option* is

```
[;SHOW[=showparmlist]]
[ONERR[OR={REDO
        QUIT}]]

[ {;DATE[S]<=accdate
  ;DATE[S]>=moddate} ]

[;PURGE]
[;PROGRESS[=minutes]]
[;DIRECTORY]

[;FILES=maxfiles]
[;TRANSPORT[=MPEXL]]
[;COPYACD][;NOACD]
[;FCRANGE=filecode/filecode[,...]]
[;MAXTAPEBUF]
[;NOTIFY]

[;STATISTICS]

[;ONVS=volumesetname[,volumesetname[,...]]]
[;SPLITVS=split_setname[,split_setname]]

[;RENAME]
[;TREE][;NOTREE]

[;INTER]
[;STORESET=(device[,...])[,(device[,...])[,...]]]

[;PART[IAL]DB]
[;FULLDB]
[;STOREDIR[ECTORY][=directoryname]]
[;NOSTOREDIRE[CTORY]]

**TurboStore/iX II options"

[;COMPRESS[=compressionparmlist]]

[;MOSET=(ldev[,...])[,(ldev[,...])[,...]]]
[;NAME=backupname]

**TurboSTORE/iX 7x24 True-Online Backup options:

[;ONLINE[={START
            END}[,time][,ASK]]]

[;LOGVOLSET=volumesetname]
```

RESTORE Syntax

```
RESTORE [ restorefile][;filesetlist][;option[;...]]
```

where option is

```
[;SHOW[=showparmlist]]
[;OWNER[OR]={QUIT
             SKIP
             FULL}]
[;{LOCAL
   GROUP=groupname
   ACC[OUN]T=accountname}[,...]]
[;CREATE[={GROUP
           ACCT
           CREATOR
           PATH}[,...]]]
[;CREATOR[=username]]
[;GID[=filegroupname]]
[;{KEEP
   NOKEEP}]
[;{OLDDATE
   NEWDATE}]
[;DIRECTORY]
[;LISTDIR]
[;PROGRESS[=minutes]]
[;FCRANGE=filecode/filecode[,...]]
[;FILES=maxfiles]
[;DEC=device]
[;VOL=volumename]
[;VOLCLASS=volclassname]
[;VOLSET=volumesetname]
[;COPYACD][;NOACD]
[;TREE][;NOTREE]
[;STOREDIR[IRECTORY]=directoryname]
[;PART[IAL]DB]
[;RESTORESET=(device[,...])[,...]]

**TurboStore/iX II and TurboSTORE/iX 7x24 True-Online
Backup options:
[;RESTORESET=(device[;...])[,(device[,...])[,...]]]
[;MOSET=(ldev[,...])[,(ldev[,...])[,...]]]
[;NAME=backupname]
```

VSTORE Syntax

```
VSTORE[vstorefile][;filesetlist][;option[;...]]
```

where option is

```
[;SHOW[=showparmlist]]  
[;ONERR[OR]={QUIT SKIP}]  
[;DIRECTORY]  
[;PROGRESS[=minutes]]  
[;COPYACD][;NOACD]  
[;TREE][;NOTREE]  
[;NODECOMPRESS]  
[;STOREDIR[ECTORY[=directoryname]]  
[;PART[IAL]DB  
[;RESTORESET=(device[,...])]
```

**TurboStore/iX II and TurboSTORE/iX 7x24 True-Online
Backup options:

```
[;RESTORESET=(device[,...])[,(device[,...])[,...]]]  
[;MOSET=(ldev[,...])[,(ldev[,...])[,...]]]  
[;NAME=backupname]
```

C Transferring Optical STORE Backups onto Standard Tapes

STORCOPY is a utility that allows you to transfer optical STORE backups onto standard tape backup, a less expensive backup medium. STORCOPY copies backups on magneto-optical devices to any other STORE backup media *except* magneto-optical.

The STORCOPY program resides in PUB.SYS. It can be used by users with OP and SM capability by entering either of the following commands from the MPE/iX prompt:

```
:STORCOPY option Return  
or
```

```
:STORCOPY Return
```

NOTE If you enter STORCOPY where no "*option*" parameters are specified, as in the second example, the "STORCOPY:" prompt appears. Acceptable responses to this prompt are a complete STORCOPY command line (see the following examples).

In the following example, a backup called MOBACK is copied to a single output device on LDEV 8. No listing is produced. MOBACK may have been created with the STORESET option, and thus may contain more than one set. STORCOPY will copy these sets one at a time to the output device until all the sets have been copied.

```
:STORCOPY COPYTOSET=(8);NAME=MOBACK Return  
or  
:STORCOPY Return  
:STORCOPY:storcopy copytoaset=(8);name=moback Return
```

The next example shows a backup called MOBACK is copied to a single output device defined by the file equation DDSTAPE to a DDS drive. A listing will be produced and sent to \$STDLIST and the line printer.

```
:FILE DDSTAPE;DEV=DDS Return  
:FILE SYSLIST;DEV=LP Return  
:STORCOPY  
COPYTOSET=( *DDSTAPE );NAME=MOBACK;SHOW=OFFLINE Return  
or
```

: STORCOPY Return

: STORCOPY:FILE DDSTAPE;DEV=DDS Return

: STORCOPY:FILE SYSLIST;DEV=LP Return

: STORCOPY:STORCOPY

COPYTOSET=(*DDSTAPE);NAME=MOBACK;SHOW=OFFLINE Return

In the following example, a backup called **MOBACK**, which has five sets, is copied to three output devices. Two of the output devices are DDS drives on LDEV 7 and LDEV 8. The other output device is defined by file equation **DDSTAPE**. A **SECURITY** and **LONG** listing is produced and sent to **\$STDLIST** and a new file called **MOLIST** via the file equation **SYSLIST**.

Since **MOBACK** consists of five sets and only three output devices, **STORCOPY** will copy the first three sets to output and will produce the other two sets as the output devices become available.

:FILE DDSTAPE;DEV=DDS Return

:FILE SYSLIST=MOLIST,NEW;SAVE Return

:STORCOPY COPYTOSET=(7),(8),(*DDSTAPE);NAME=MOBACK; & Return

:SHOW=SECURITY,OFFLINE Return

or

: STORCOPY Return

: STORCOPY:FILE DDSTAPE;DEV=DDS Return

: STORCOPY:FILE SYSLIST=MOLIST,NEW;SAVE Return

: STORCOPY:STORCOPY COPYTOSET=(7),(8),(*DDSTAPE);NAME=MOBACK; Return

:SHOW=SECURITY,OFFLINE Return

NOTE Parenthesis () enclosed in quotation marks *must* be entered as part of the command line syntax.

COMMAND `(parameter)`

Glossary

absolute pathname A pathname that begins with the root directory, such as `/SYS/PUB/TDP`. See also *pathname* and *relative pathname*.

access control definition (ACD) Security feature that controls access to files and directories. Consists of a list of access permissions and user specifications. (For example, `R,W,X:@.PAYROLL` gives all users in the `PAYROLL` account read, write, and execute access to the file or directory that is assigned this ACD.) ACDs are applied to files or directories by using the `ALTSEC` command. By default, all files existing outside the traditional MPE account/group structure and all directories are assigned ACDs when they are created.

access mode A type of access permitted to a file, such as write, read, or execute access.

appropriate privilege Having sufficient capabilities to perform an operation on MPE/iX. SM capability always provides appropriate privilege to system administrators.

archiving Storing infrequently used or out-of-date files onto tape and permanently removing them from the system disks at the same time.

back reference The technique of using an asterisk (*) before a formal file designator to indicate that it has been previously defined during the current session or job with the `FILE` command.

backup The process that duplicates computer data to offline media, such as a magnetic tape. Backups protect data if a system problem should occur.

backup devices System peripherals that allow you to write information to, and read information from, backup media.

case sensitivity HFS filenames can be saved in uppercase or lowercase letters. The file named `./FILE1` does not refer to the same file as `./file1` or `./File1`.

character A letter, number, or symbol represented by one byte of data.

character set A series of characters to substitute for a single character to name a group of files.

command A system-reserved word that directs the operating system, a subsystem, or a utility program to perform a specific operation.

compatibility mode(CM) Compatibility mode provides object code compatibility between MPE V/E-based systems and the 900 Series HP 3000. Compatibility mode allows current Hewlett-Packard customers to move applications and data to the 900 Series HP 3000 without changes or recompilation.

concurrent backup devices A set of multiple, concurrently accessible backup devices to which you are able to store information simultaneously.

consecutive backup devices A set of multiple backup devices used for a given tapeset to which you are able to store information sequentially.

console See system console.

current working directory The directory in which you are working and from which relative pathnames are resolved. See also *directory* and *relative pathname*.

DAT Digital Audio Tape. A data storage media used by the HP 1300H DDS-format device.

database An integrated collection of logically related data files and the structural information about the data.

data compression The use of data compression algorithms on the host system to reduce the amount of data transferred to the output device.

DDS Digital Data Storage. A format for storing computer data on DAT cassettes.

device See peripheral.

device class A collection of devices that have some user-defined relation. The MPE/iX file system supports a means of maintaining collections of devices but is cognizant of the relation only for certain MPE device class names.

device pool A series of identical backup devices.

device type Device types are defined by a number. For example, 0 represents a moving-head disk and 16 is the class type number assigned to line printers.

directory A special kind of file that contains entries that point to other files. It acts like a container for files and other directories. On MPE/iX, accounts and groups are special types of directories.

disk A circular plate of magnetically coated material used to store computer data. A disk may be fixed, removable, hard, or flexible.

dot (.) Convention that signifies the current working directory in HFS syntax. See also *current working directory* and *HFS syntax*.

dot-dot (..) Convention that signifies the parent directory in HFS syntax. See also *current working directory* and *HFS syntax*.

dot-slash (./) A convention that you can use to refer to a file using HFS syntax in your current working directory. For example, `./myfile.old` refers to the file `myfile.old` in your current directory.

error messages Messages describing errors occurring during either an interactive session or a batch job. The messages are reported to the standard list device, which is usually a terminal (for a session) or a line printer (for a job).

file An object that can be written to, read from, or both. A file has certain attributes including access permissions and file type. A group of related records that represents ASCII text (text files) or binary data (such as executable code). Every file must have a filename so that the user can access the file's contents.

file equation A method of equating a filename to a device or another file. The MPE/iX `FILE` command is used to establish the relationship of the file to the device. Generally used to direct the input to or output from a program, job, or session to a particular device by referencing the device class, such as `TAPE` or `LP`.

filename A name of a file that can be in MPE syntax (`FILE.GRP.ACCT`) or HFS syntax (`/ACCT/GRP/FILE1`). Each syntax has different restrictions on filename length and the characters that can compose the name. See also *MPE syntax* and *HFS syntax*.

file owner The person who has complete access to a file unless the user is restricted by a `$OWNER ACD` entry. The `$OWNER ACD` entry can restrict the file access of the file owner. The file owner is similar to the file creator. The command `LISTFILE filename -3` displays the fully qualified user ID (`user.account`) of the file owner.

group For POSIX compatibility, refers to a group of related users. This is distinct from MPE groups, which are special types of directories existing directly below accounts.

group ID database A system database that contains the group name, group ID, and user names for all groups.

group ID (GID) A number that determines group access privileges. (On MPE/iX, it is actually the string `@.account`).

HFS syntax Expanded MPE/iX syntax that is case sensitive and allows users to address multiple levels in the hierarchical file system. A name beginning with `./` or `/` automatically signifies HFS syntax to MPE/iX.

Some additional rules are as follows:

- Names of directories directly under root or under an account may have up to 16 characters.
- Names of directories or files not directly under the root or a group can be up to 255 characters.
- Names of directories and files can contain the following special characters: hyphen (-), dot (.), or slash (/)

hierarchical file system (HFS) A file system that is tree structured and can contain files at many different levels. This file organization is obtained through the use of directories, which can contain files and other directories.

indirect file A text file containing the parameters for a `STORE` or `RESTORE` command that you execute regularly.

input The data to be processed, or the process of transferring data from external storage to the computer.

input/output (I/O) The process of, or equipment used in, transmitting information to or from the computer.

interleave A method of writing data to a disk device for the purpose of improving data access speed.

I/O error A data transmission error between a computer and peripheral. Examples of I/O errors are baud rate or parity mismatch, and incorrect syntax in device-control instructions.

job A single file, submitted by a user, containing operating system and utility commands and references to the files to be manipulated.

media Devices capable of storing data, such as disks or magnetic tape.

media name The name of the media set to make available for use by TurboSTORE/XL II. The media name is the name assigned during the `STORE` operation.

media sub-name The name of the surface within the media set to use. The media sub-name is the sub-name assigned during the `STORE` operation.

mounting The act of making a data storage device accessible. To physically mount the device, you load the media onto the device. To logically mount the device, you tell the operating system which device you want to use, and it allows you access to that device.

MPE/iX Multi-programming executive with integrated Posix: The operating system for 900 Series HP 3000 computers. MPE/iX manages all system resources and coordinates the execution of all programs running on the system.

MPE syntax Rules that determine the filename length, special characters, and conditions for files, groups, and accounts. Account, group, and filenames can be up to 8 characters. Characters are always converted up uppercase. Characters must be alphanumeric. This is the syntax current MPE/iX users are used to using (for example, `LEDGER.PUB.SYS`).

native mode (NM) The native run-time environment of MPE/iX. In native mode, source code has been compiled into the native instruction set of the 900 Series HP 3000.

online backup Online backup means that the `STORE` fileset and structures are accessible for any access while the backup is taking place. Any modifications made to the `STORE` fileset during the backup are logged and saved along with the data on the backup medium. On `RESTORE`, the data and log file are used to recover the data to a consistent state.

output Data transferred from internal to secondary (external) storage in a computer. Also, the process of transferring information from the computer to a peripheral device.

parallel devices A set of backup devices to which you are able to store information simultaneously.

parameter A value in a list of values that is passed to a procedure. The parameter is used on calculations or operations in the procedure.

parent directory A directory that contains other directories.

pathname A way of identifying the path to any MPE/iX file. For example, you can refer to FILE1.PUB.SYS using the pathname /SYS/PUB/FILE1. Notice that pathnames are top-down rather than bottom-up as MPE syntax.

peripheral A hardware device attached to and controlled by a computer, such as a terminal, a tape or disk drive, or a printer.

POSIX Portable Operating System Interface. A set of standards that address various areas of operating system technology. The POSIX standards describe functions of an operating system interface that applications use to become "POSIX-compliant." The main point of POSIX is to facilitate software portability and minimize porting costs.

purge To delete a permanent file from the system with the PURGE command. The PURGE command is also used to delete an account structure entry such as a user name, a group name, or an account.

relative pathname A pathname that is interpreted from the current working directory. For example, ./dir1/longfilename refers to the file longfilename in directory dir1 in the current working directory.

required parameter A parameter that is required when entering a command or calling an intrinsic. In reference manuals, required parameters are surrounded by braces ({}).

restore The process of retrieving user files from SYSGEN (MPE/iX), SYSDUMP (MPE V/E), and STORE tapes or serial disk and writing them to disk. Restoring is executed with the RESTORE command.

root directory Also called (and designated by) a slash or /. It is a system directory; all files, accounts, groups, and directories connect back to the root directory. All accounts on MPE/iX are direct descendants of the root directory.

scratch tapes Used tapes containing information that is no longer needed.

sequential A manner in which information may be read from or written to a device. Sequentially accessed files are stored in such a way that the logical order of the file's record is identical to their physical layout on disk or tape.

shadow logging A method of saving the "before" images of file blocks when the records in the file block are modified during online backup. The before images are stored on tape and are used by RESTORE to reconstruct the original contents of the file.

slash (/) Another name for the root directory. See also *root directory*.

storage device A device (such as a disk pack, a disk cartridge, a flexible disk, magnetic tape, or cartridge tape), onto which data can be stored and subsequently retrieved.

subdirectory A directory that is contained within another directory is sometimes referred to as a subdirectory.

system A group of one or more CPUs that communicate through buses without the use of data communications software.

sync point The point in time the backup occurs, where all files are synchronized. It can occur at different times, depending on the type of backup performed. You can schedule when the sync point occurs by scheduling when the backup job is launched.

system console 1)A workstation given a unique status by the operating system. It is used by the operator to execute specific commands for the purpose of managing sessions, jobs, and system resources. It is associated with all boot or system loader error messages, system error messages, and certain system status messages.

2)The terminal, usually logical device 20, that the system operator uses to monitor system activity, respond to resource requests, and send messages to user's terminals. The console (and its associated privileges and responsibilities) may be transferred to another logical device with the `CONSOLE` command.

system file directory A directory maintained by MPE that records the name, group, and account of each permanent file on the system. The directory contains the size of each file, its location on the disk, who may access it, and other information.

system manager The person who manages the computer installation, responsible for creating accounts and assigning capabilities and resource-use limits to each.

tape request A printed message at the console asking for a backup device to be assigned to a user.

transport backup Transport refers to the movement of files between MPE V/E and MPE/iX systems using backup media. MPE/iX backup and recovery provides a compatibility option with which an MPE V/E-compatible tape may be created or read.

true online backup A tool that allows users to perform application and system backups without closing files or disrupting users.

user Anyone logged on to a session, using a local or remote terminal to interact with the computer. Each user is identified by a user and account name, and can access files in the logon group.

utility program An operating system program that performs specific functions such as file copying, sorting and merging, memory dump analysis, or monitoring available disk space.

volume A volume is one disk. Each volume is a member of a volume set and contains a volume label, a label table, and a free space map.

volume class Volume classes are used for the allocation and restriction of disk space. A volume class is a logical subset or partition within a volume set and can bridge any number of physical member volumes of a volume set.

volume set A volume set is a group of from 1 to 255 related disks. One volume of the volume set must be designated as the master volume for the set. Each volume set is assigned a name by which it is identified and referenced.

wildcard A symbol that is used to replace a character or set of characters. In MPE, the "at sign" (@), the "pound sign" (#), and the question mark (?) are used as wildcard characters.

write-enable To remove a disk's write-protection, allowing the disk to be written upon.

write-protect To protect stored data so that it cannot be overwritten.

write ring A plastic ring that fits onto the inner groove of a reel-to-reel tape, enabling you to write information onto the tape.

Symbols

\$NULL, 52
\$STDIN, 151
\$STDLIST, 151

Numerics

7x24 true-online backup, 25, 114, 119

A

aborting sessions, 64
ABORTJOB command, 64
account, 171
ACCOUNT parameter, RESTORE command, 172
account structure, 92
ALLBASE/SQL databases, storing, 73
allocating a device, 135
allowing users on system, 146
archiving files, 91
ASK parameter
 ONLINE option, 128
assigning an LDEV number, 135, 137
attach phase, 25, 115, 116
automatic volume recognition, 75
automatically assigning devices, 46, 47, 135, 139
AVR, 75

B

back up
 files, 35
backing up files, 117
backup
 7x24 true-online, 25, 114
 error recovery, 78, 89
 I/O error, 78, 89
 log files, 117
 multiple device, 47
 of account structure, 92
 online, 25
 parallel, 49
 parallel device pools, 51
 record size, 96
 request, 38, 77
 sequential, 48
 single-device, 46
backup companies, commercial, 189
backup devices
 preparing, 46
 supported by MPE/iX, 20
backup error recovery, 78
backup media

 defined, 20
 duplicate, 191
 preparing, 40
 verifying, 143
backups
 labeled, 37
 unlabeled, 37
base group, SYSGEN utility, 150
basegroup parameter, SYSGEN utility, 151
BREAKJOB command, 62
buffer size, 96

C

caret, 71
character sets, 66
cold site, 188
commercial backup companies, 189
compatibility
 of MPE V/E & XL backups, 93
 with MPE V/E, 24, 37
COMPRESS parameter
 HIGH option, 113
 LOW option, 113
 TurboSTORE/iX II command, 113
compressing data, 113
compression algorithms, 113
computer service bureaus, 189
CONFIG.SYS, 150, 151
configurators, SYSGEN utility, 150, 153
console
 checking the, 136
 requests, 61
CREATE parameter, RESTORE command, 171, 174
creating labeled volumes, 75
creator, 171
CREATOR parameter, RESTORE command, 172, 174

D

DASS, 24
data compression, 113
databases
 full backup, 109
 partial backup, 109
 quiescing, 123
 restoring part of, 178
 storing ALLBASE/SQL, 73
 storing non-HP, 74
 storing TurboIMAGE, 73
DATE parameter

- state change date, 90
- STORE command, 90
- default group, SYSGEN utility, 150
- DEV parameter
 - STORE command, 176, 177
- device class, 46
- device pool, 39, 48, 49
- device, restoring to specific, 176
- directory information, restoring, 174
- DIRECTORY parameter
 - RESTORE command, 174, 184
 - STORE command, 92
 - VSTORE command, 143
- disaster recovery, 187, 192
- disaster recovery plan, 21
- disk read errors, 89
- displaying
 - file dates, 86
 - file security information, 87
 - progress messages, 92, 175
 - tape information, 174
- duplicate backups, 191
- duplicate filenames, restoring to, 159

E

- END parameter
 - ONLINE option, 128
- error check for STORE, 143
- error recovery
 - ONERROR option, 89
 - RESTORE command, 170
 - STORE, 89
- errors
 - determining the cause of, 194
 - recommended recovery procedures, 195
- exclamation point, 71

F

- FCRANGE parameter
 - RESTORE command, 176
 - STORE command, 96
- file codes, 96, 176
- FILE command, 38, 46, 65, 77, 139, 180
 - automatically assigning devices, 167
 - restoring files, 161, 162, 167
- file dates, displaying, 86
- file equation, 75
- file security information
 - displaying, 87
- file system directory, 29, 149, 183
- FILES= parameter

- RESTORE command, 176
- STORE command, 96
- fileset, 65, 160
- forms requests, 61
- full system backup, 30, 91, 112, 183
- FULLDB option
 - STORE command, 109

G

- group, 171
- GROUP parameter, RESTORE command, 172, 174

H

- hardware requirements, 18
- HELLO command, 59
- HIGH option
 - COMPRESS parameter, 113
- HIPRI parameter
 - HELLO command, 59
 - JOB command, 59
- hot site, 188

I

- idx aborts, 194
- improving storage efficiency, 96, 113
- increasing record size, 96
- indirect file, 70, 89
 - restoring files, 170
- initializing volumes, 184
- inputfile parameter, SYSGEN utility, 151
- INTER parameter
 - STORE command, 112
- interleave format, 112

J

- JOB command, 59
- job priority, 59
- JOBFENCE command, 59, 146
- jobs, 59

K

- KEEP command
 - SYSGEN utility, 150, 151
- KEEP parameter
 - RESTORE command, 159, 173, 184

L

- labeled backups, 37, 75
- labeled tapes, 24

- restoring files, 161, 170
- labeled volumes, 75
- last access date, 173
- LDEV, 36
- LIMIT command, 59, 146
- LISTDIR parameter
 - RESTORE command, 174
- LISTF command, 159
- listing the files restored, 169
- listing the files stored, 83, 87
- LOCAL parameter
 - RESTORE command, 172, 174
 - VSTORE command, 143
- log files, 117
- log handler, 25, 115, 116
- logging off users, 57
- logical device (LDEV) number, 36, 37, 46, 77, 135, 137, 176, 180
- LOGVOLSET option
 - TurboSTORE/iX 7x24 True-Online Backup command, 117
- long file display, 83, 86
- LOW option
 - COMPRESS parameter, 113

M

- maintaining a backup library, 32
- maximum number of files, restoring, 176
- maximum number of files, specifying, 96
- MAXTAPEBUF parameter
 - STORE command, 96
- media error recovery, 78, 89
- media request, 61
 - automatic answering, 135
 - denying, 138
 - handling, 135, 138
 - pending, 138
 - replying to, 137
- modification date, 173
- monitoring RESTORE command progress, 180
- monitoring STORE command progress, 77
- mounting additional backups, 79
- mounting additional tapes, 180
- MPE V/E
 - compatibility, 24, 37
 - compatible backups, 93
 - tape format, 31
- multiple disk files
 - restoring files, 165
- multiple-device backup, 31, 35, 37, 38, 65, 77
 - described, 47
 - planning use of, 37

- multiple-device planning chart, 36
- multiple-device RESTORE, 180
- mutual backup arrangements, 188

N

- naming files
 - in RESTORE command, 160
 - in STORE command, 66, 70
- native mode, 83
- NEWDATE option
 - RESTORE command, 173
- newgroup parameter, SYSGEN utility, 151
- NEWVOL command
 - VOLUTIL utility, 184
- NO WRITE RING message, 79
- NOKEEP option
 - RESTORE command, 173
- nonsystem volume set directory, 92
- NOSTOREDIRECTORY option
 - STORE command, 110
- NOTIFY parameter
 - STORE command, 97
- notifying users of a backup, 58

O

- OLDDATE option
 - RESTORE command, 173
- ONERROR option
 - error recovery, 89
 - REDO parameter, 78, 89
 - RESTORE command, 170
 - STORE command, 78, 89
 - VSTORE command, 143
- ONERROR=QUIT option
 - VSTORE command, 143
- ONERROR=SKIP option
 - VSTORE command, 143
- online backup, 25, 114
- ONLINE option
 - ASK parameter, 128
 - END parameter, 128
 - specifying time, 128
 - START parameter, 128
 - STORE command, 116
- ONVS parameter
 - STORE command, 92, 103
- outputfile parameter, SYSGEN utility, 151

P

- parallel device backup, 49
 - preparing backups, 45

- parallel device pools, 39, 51, 77, 112
 - preparing backups, 45
 - restoring devices, 163
- parallel devices, 39, 49, 50, 77, 78, 89
 - maximum number of, 50
 - restoring files, 163, 164
- parallel disk files
 - restoring files, 165
- partial system backup, 30, 90, 112, 183
- PARTIALDB option, 131
 - restoring part of a database, 178
 - STORE command, 109
- placeholder, 39, 49, 66
 - restoring files, 162, 167
- preparing backup media, 40
- preparing storage devices, 35
 - LDEV number, 36
 - multiple-device backup, 35
 - multiple-device planning chart, 36
 - selecting backup method, 37
- preparing the system for backup, 57
- preventing new system activity, 59
- printing the files restored, 169
- printing the files stored, 83
- private recovery facility, 188
- process identification number (PIN), 135, 137
- program aborts, 194
- progress messages, 79
 - displaying, 92, 175
- PROGRESS parameter
 - RESTORE command, 175, 181
 - STORE command, 79, 92
- PUB group, 29
- PURGE parameter
 - STORE command, 91
- Q**
- quiescing databases, 123
- QUIET mode, 58, 60
- QUIT parameter; ONERROR option, 78, 89, 170
- R**
- read access, 172
- RECALL command, 61, 136, 138
 - CTRL/A keys, 138
- record size, 96
- recovery facility, 188, 192
- recovery facility response time, 190
- recovery from backup errors, 78, 89
- recovery from tape errors, 170
- recovery system compatibility, 190
- redirecting list of files stored, 83
- REDO parameter
 - ONERROR option, 78, 89
- reel-to-reel media
 - MPE/iX backup media, 20
- reel-to-reel tapes
 - end-of-file marker, 31
 - header label, 31
 - labeled, 31
 - unlabeled, 31
- reload of system configuration, 183, 184
- removing files from disk, 91
- REPLY command, 137, 138
 - CTRL/A keys, 138
- replying to console requests, 61
- RESTORE
 - monitoring progress, 180
 - using PARTIALDB, 178
 - using RESTORESET, 162
 - using STOREDIRECTORY, 178
 - using VSTORE, 166
- RESTORE command, 26, 96, 112, 157, 183, 185
 - ACCOUNT parameter, 172, 174
 - CREATE parameter, 160, 171, 174
 - CREATOR parameter, 172, 174
 - DIRECTORY parameter, 174, 184
 - error recovery, 170
 - FCRANGE parameter, 176
 - FILE parameter, 167
 - FILES=parameter, 176
 - GROUP parameter, 172, 174
 - KEEP parameter, 159, 173, 184
 - LISTDIR parameter, 174
 - LOCAL parameter, 172, 174
 - NEWDATE parameter, 173
 - NOKEEP parameter, 173
 - OLDDATE parameter, 173
 - ONERROR parameter, 170
 - options, 169
 - placeholder, 167
 - preparing to restore files, 159
 - PROGRESS parameter, 175, 181
 - referencing files, 160
 - RESTORESET parameter, 162, 163
 - SHOW parameter, 169, 182
 - SHOW=OFFLINE parameter, 182, 186
 - syntax, 199
 - TRANSPORT parameter, 176
 - unrecoverable errors, 170
- RESTORESET parameter
 - RESTORE command, 162, 163
 - STORE command, 70

-
- restoring
 - directory information, 174
 - part of a database, 178
 - using a disk directory file, 178
 - restoring files, 20
 - 7x24 true-online backups, 132
 - account, 171
 - automatically assigning devices, 167
 - checking for duplicate filenames, 159
 - creator, 171
 - described, 157
 - error recovery, 170
 - FILE command, 162
 - from a disk file, 164
 - from labeled tapes, 161
 - from multiple disk files, 165
 - from parallel device pools, 163
 - from parallel devices, 163
 - from parallel disk files, 165
 - from sequential devices, 162
 - from single devices, 161
 - full backups, 186
 - group, 171
 - issuing a RESTORE command, 167
 - LISTF command, 159
 - listing the files restored, 169
 - maximum number, 176
 - overwriting files, 173
 - parallel devices, 164
 - partial backups, 185
 - placeholder, 162
 - preparation, 159
 - printing the files restored, 169
 - read access, 172
 - referencing files, 160
 - sequential devices, 163
 - STOREDIRECTORY option, 132
 - tape error recovery, 170
 - to specific device, 176
 - to specific volumes, 176
 - using VSTORE, 166
 - wildcard characters, 159
 - write ring, 159
 - restoring from
 - compressed files, 113
 - multiple devices, 162
 - restoring TRANSPORT backup to MPE V/E, 96
 - RESUMEJOB command, 146
 - returning system to normal, 146
 - scratch tapes, 32
 - sequential backup, 48
 - sequential device backup
 - preparing backups, 40, 44
 - sequential devices, 39, 48, 49, 77, 78, 89
 - maximum number of, 48
 - restoring files, 162, 163
 - service bureaus, computer, 189
 - sessions, 59
 - shadow logging, 125
 - shadow subsystem, 25, 115, 116
 - short file display, 83, 86
 - SHOW parameter
 - RESTORE command, 169, 182
 - STORE command, 83, 87, 142
 - VSTORE command, 143
 - SHOW=DATES option
 - STORE command, 86
 - SHOW=LONG option
 - STORE command, 85
 - SHOW=OFFLINE option
 - RESTORE command, 182, 186
 - STORE command, 32, 83, 145
 - SHOW=SECURITY option
 - STORE command, 87
 - SHOW=SHORT option
 - STORE command, 84
 - SHOWJOB command, 59, 60, 62, 137, 146
 - SHOWVAR, 194
 - single-device backup
 - described, 31, 46
 - preparing, 40
 - preparing backups, 44
 - SKIP parameter;ONERROR option, 170
 - specifying maximum number of files, 96
 - specifying time, 128
 - START command, initial system loader (ISL), 154
 - START parameter
 - ONLINE option, 128
 - state change date
 - DATE parameter, 90
 - STATISTICS parameter
 - STORE command, 98
 - stopping jobs, 62
 - storage devices
 - preparing, 35
 - STORE
 - error recovery, 89
 - unrecoverable errors, 89
 - STORE command, 26, 65
 - DATE parameter, 30, 90

- DEV parameter, 176, 177
- DIRECTORY parameter, 92
- FCRANGE parameter, 96
- FILES= parameter, 96
- FULLDB option, 109
- INTER parameter, 112
- MAXTAPEBUF parameter, 96
- monitoring progress, 77
- NOSTOREDIRECTORY, 110
- NOTIFY parameter, 97
- ONERROR option, 78, 89
- ONLINE option, 116
- ONVS parameter, 92, 103
- options, 81
- options (parameters), 65
- PARTIALDB option, 109
- PROGRESS parameter, 79, 92
- PURGE parameter, 91
- referencing files, 66, 70
- RESTORESET parameter, 70
- SHOW parameter, 83, 87, 142
- SHOW=DATES option, 86
- SHOW=LONG option, 85
- SHOW=OFFLINE option, 32, 83, 145
- SHOW=SECURITY option, 87
- SHOW=SHORT option, 84
- STATISTICS parameter, 98
- STOREDIRECTIONS, 110
- storefile parameter, 52
- STORESET parameter, 39, 49, 70, 97, 112
- syntax, 198
- TRANSPORT parameter, 24, 31, 37, 84, 85, 93
- VOL parameter, 176
- VOLCLASS parameter, 176
- VOLSET parameter, 176
- STORE media, 29
- STORE program
 - described, 29
- STOREDIRECTIONS option
 - restoring files, 132
 - restoring using a disk directory file, 178
 - STORE command, 110
- storefile parameter, 52
- STOREJCW, 194
- STORESET parameter
 - STORE command, 39, 49, 70, 97, 112
- storing compressed files, 113
- storing efficiently, 96, 113
- storing files, 20, 65
 - by date, 91
- storing online, 25, 114
- storing to
 - a single device, 46
 - labeled backups, 75
 - multiple devices, 47
 - parallel device pools, 51
 - parallel devices, 49
 - sequential devices, 48
- streaming backup drive, 112
- suspended job, restarting, 146
- suspending jobs, 62
- sync point, 120
 - beginning of backup, 121
 - end of backup, 121
- SYS account, 29, 150, 151
- SYSGEN utility, 149, 154
 - base group, 150
 - basegroup parameter, 151
 - default group, 150
 - described, 29
 - KEEP command, 150
- SYSGIN, 151
- SYSGOUT, 151
- SYSLIST, 78, 83, 86, 89
- system backup
 - aborting sessions, 64
 - defined, 16
 - full, 30, 91, 112
 - log files, 117
 - logging off users, 57
 - multiple-device, 31, 47, 65
 - notifying users, 58
 - online, 114
 - partial, 30, 90, 112
 - preparing the system, 57
 - scheduling, 30
 - sending a WARN message, 60
 - single-device, 31, 46
 - suspending jobs, 62
- system compatibility for recovery, 190
- system configuration, 29, 183
 - backup, 149
 - reload, 183, 184
- system directory
 - backup, 92
- system failure, 16, 183, 187
- system files, 29, 149, 183
- system load backup, 149, 183, 184
 - creating, 152, 154
- system load media, 29
- system reload, 183
- system volume set directory, 92

T

TAPE command, SYSGEN utility, 152
tape error recovery, 170
tape information, displaying, 174
tape label, 40
tape request, 180
tapes
 labeled, 24
 unlabeled, 24
task roadmap, 26
TELL command, 58, 138
testing recovery system, 192
transport mode, 84
TRANSPORT parameter
 backup, restoring to MPE V/E, 96
 RESTORE command, 176
 STORE command, 24, 31, 37, 84, 93
true-online backup, 25
TurboIMAGE databases, storing, 73
TurboSTORE/iX 7x24 True-Online Backup
 described, 119
 LOGVOLSET option, 117
 ONLINE option, 114, 128
 options, 82, 128
 PARTIALDB option, 131
 quiescing databases, 123
 restoring, 132
 shadow logging, 125
 specifying time, 128
 sync point, 120
TurboSTORE/iX II command
 COMPRESS parameter, 113
 options, 82

U

unlabeled backups, 37
unlabeled tapes, 24
 restoring files, 170
unlabeled volumes, 75
unrecoverable errors
 RESTORE command, 170
 STORE, 89
user files, 29, 149, 183
user task roadmap, 26
users, helping to store & restore, 26

V

verifying backup media, 143
VOL parameter
 STORE command, 176
VOLCLASS parameter

 STORE command, 176
VOLSET parameter
 STORE command, 176
volume set, 103
volumes, 176
VOLUTIL utility, 184
 NEWVOL command, 184
VSTORE command, 65, 143, 144
 ONERROR=QUIT option, 143
 ONERROR=SKIP option, 143
 restoring files, 166
 syntax, 200

W

WAIT state, 59
WARN command, 60
WELCOME message, 57
wildcard characters, 66, 70
 restoring files, 159
write ring, 79
 restoring files, 159