



RTE-A Backup and Disk Formatting Utilities

Reference Manual

**Software Services and Technology Division
11000 Wolfe Road
Cupertino, CA 95014-9804**

NOTICE

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

HEWLETT-PACKARD MAKES NO WARRANTY OF ANY KIND WITH REGARD TO THE MATERIAL, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, 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 are reserved. No part of this document may be photocopied, reproduced, or translated to another language without the prior written consent of Hewlett-Packard Company.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the 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.

Copyright © 1993, 1995 by Hewlett-Packard Company

Printing History

The Printing History below identifies the edition of this manual and any updates that are included. Periodically, update packages are distributed which contain replacement pages to be merged into the manual, including an updated copy of this printing history page. Also, the update may contain write-in instructions.

Each reprinting of this manual will incorporate all past updates; however, no new information will be added. Thus, the reprinted copy will be identical in content to prior printings of the same edition with its user-inserted update information. New editions of this manual will contain new information, as well as all updates.

To determine what manual edition and update is compatible with your current software revision code, refer to the Manual Numbering File. (The Manual Numbering File is included with your software. It consists of an "M" followed by a five digit product number.)

First Edition	Nov 1993	Rev. 6100 (Software Update 6.1)
Second Edition	Apr 1995	Rev. 6200 (Software Update 6.2)

Preface

This manual describes the scope, format, and use of backup and disk formatting utilities available on RTE-A. For a complete list of the RTE-A manual set, refer to the *RTE-A Index and Glossary*, part number 92077-90036. This manual is organized as follows:

- Chapter 1 Provides an overview of the RTE-A backup and disk formatting utilities.
- Chapter 2 Contains a detailed description of each of the utilities mentioned in Chapter 1. They are organized alphabetically.

Conventions Used in This Manual

The command syntax and other conventions used in this manual are described in the following paragraphs. Sample terminal displays include both user input and program prompts and messages. Comments are given in parentheses. For example, entering a command at the CI> prompt is shown as:

```
CI> dl /derick/casey/@@ (List all files in subdirectory CASEY
                        under directory DERICK)
```

The command syntax conventions are as follows.

Convention	Meaning
Uppercase letters	Capital letters indicate the exact characters required. However, CI accepts lowercase input. For example, the command syntax for the AS command is: <pre>AS prog <partition number> [C D]</pre> and the actual sample entry can be: <pre>CI> as testprogram 2 c</pre>
[]	Optional parameters are shown in square brackets. If you omit a parameter, use a comma as a placeholder before specifying additional parameters.
	The vertical bar separates alternate choices. The “C D” shown above indicates either C or D can be entered.
, or blank	Use a comma or a space as a delimiter between commands and parameters. Blank spaces are used throughout this manual in all syntax strings.
<i>lowercase italics</i> and < >	Lowercase italicized letters represent user-supplied variables. In long descriptive phrases, the variable may be enclosed in angle brackets for clarity, for example, <partition number>.

Table of Contents

Chapter 1 Introduction

Backup and File Interchange Utilities	1-1
Backup Utilities	1-2
File (Logical) Backup	1-2
Physical Backup	1-3
Compatibility with Other Disk Backup Utilities	1-3
Compatibility Among Disks	1-3
File Interchange on RTE-A	1-4
Formatting Utilities	1-5
General Formatting Information	1-5
The Preamble	1-7
The Postamble	1-7
Initializing and Sparing Hard Disks	1-7
Sparing vs. Skipping	1-8
Interleave Factor Calculation	1-8
CS/80 Cartridge Tape Drive Definitions	1-9

Chapter 2 Disk Backup and Formatting Utilities

ARSTR	2-2
Calling ARSTR	2-2
ARSTR Commands	2-2
Abort, End, and Exit (AB) (EN) (EX)	2-3
Help (HE)	2-4
List Header (LH)	2-4
List Device (LL)	2-4
Restore (RE)	2-5
Unit Restore:	2-6
RE Command Options	2-6
Rewind (RW)	2-7
Tape LU (TA)	2-7
User Error Handling (UE)	2-8
ARSTR Command Examples	2-8
The Restore Operation	2-9
Checking the Disks	2-9
Sectors per Track Must Match on Disks	2-9
Total Number of Tracks May Differ	2-9
Disk Locking	2-10
Order of Disk Restoration	2-10
Unit Restore	2-10
Save Definition Records	2-10
Restoring the Data	2-11
Verification	2-11
Break Detection	2-12
Loading ARSTR	2-12
Offline System	2-12

ARSTR Example	2-13
ASAVE and ARSTR Error Messages and Warnings	2-14
ASAVE Messages	2-14
ASAVE and ARSTR Shared Error Messages	2-15
ASAVE and ARSTR Shared Warning Messages	2-17
ARSTR Error Messages	2-18
ARSTR Warning Messages	2-18
ASAVE	2-19
Calling ASAVE	2-19
ASAVE Commands	2-20
Abort, End, and Exit (AB) (EN) (EX)	2-21
Help (HE)	2-21
List Header (LH)	2-21
List Device (LL)	2-22
Rewind (RW)	2-22
Save (SA)	2-22
SA Command Options	2-23
Disk-to-tape Verify (VE)	2-24
Tape LU (TA)	2-24
Title (TI)	2-24
User Error Handling (UE)	2-25
ASAVE Command Examples	2-26
The Save Operation	2-26
Unit Save LU Checking	2-26
Disk Locking	2-26
Tape Positioning	2-27
Save File Records	2-27
Saving the Disks to Tape	2-27
Verification	2-28
Break Detection	2-29
Loading ASAVE	2-29
ASAVE Tape and File Formats	2-29
ASAVE Tape Records	2-30
Header Records	2-30
Save Definition Record	2-31
Data Record	2-31
End-of-Data Record	2-32
ASAVE Example	2-32
ASAVE Error Messages	2-33
COPYL	2-34
Calling COPYL	2-34
The Copy Operation	2-35
COPYL Error Messages	2-35
DRSTR	2-36
Calling DRSTR	2-36
DRSTR Commands	2-36
Abort, End, and Exit (AB) (EN) (EX)	2-37
Help (HE)	2-37
List Header (LH)	2-38
List Device (LL)	2-38
Media (ME)	2-39
Restore (RE)	2-39
RE Command Options	2-41
User Error Handling (UE)	2-41

The Restore Operation	2-42
Checking the Disks	2-42
Sectors per Track Must Match on Disks	2-42
Total Number of Tracks May Differ	2-42
Disk Locking	2-43
Order of Disk Restoration	2-43
Unit Restore	2-44
Restoring Fixed Disks from Removable Disks	2-44
Header Records	2-45
Restoring the Data	2-45
Verification	2-46
Break Detection	2-46
Loading DRSTR	2-46
DSAVE and DRSTR Error Messages	2-47
DSAVE Error Messages	2-47
DSAVE and DRSTR Shared Error Messages	2-48
DRSTR Error Messages	2-50
DSAVE	2-52
Calling DSAVE	2-52
DSAVE Commands	2-53
Abort, End, and Exit (AB) (EN) (EX)	2-54
Help (HE)	2-54
List Header (LH)	2-54
List Device (LL)	2-56
Media (ME)	2-56
Save (SA)	2-57
SA Command Options	2-57
Title (TI)	2-58
User Error Handling (UE)	2-59
The Save Operation	2-59
Unit Save LU Checking	2-60
Disk Locking	2-60
Removable Disk Overwriting	2-60
Save File Records	2-61
Saving to the Removable Disk	2-61
Verification	2-62
Break Detection	2-62
Loading DSAVE	2-63
DSAVE Removable Disk Media and File Formats	2-64
Removable Disk Media Format	2-64
Removable Disk Media Records	2-64
Header Records	2-64
DSAVE Error Messages	2-66
ERTSH	2-67
ERTSH Structure	2-67
ERTSH Commands	2-68
Help (??)	2-69
Build Test Pattern (BP)	2-69
Define Cylinder Range (CY)	2-69
EXIT (EX)	2-70
Define Disk Head Range (HD)	2-70
Extra Log (LL)	2-70
Selecting Disk LU to Test (LU)	2-71

Set Test Pass Count (P#)	2-71
Execute Read-Only Error Rate Test (RO)	2-72
Define Disk Sector Range (SC)	2-72
Spare Sectors (SS)	2-73
Select Test Pattern (TP)	2-73
Execute Write/Read Error Rate Test (WR)	2-74
Using ERTSH	2-75
Disks Formatted with FORMF	2-75
Initializing and Testing a New Disk	2-75
Typical Use of ERTSH	2-75
Break Detection	2-78
Specialized Use of ERTSH	2-78
Installation Procedure Summary	2-79
Error Interpretation	2-80
Error Messages	2-83
FC	2-85
Calling FC	2-85
FC Commands	2-86
Command Summary Function (?)	2-88
Abort (AB)	2-89
Name Comment File (CF)	2-89
Cartridge List (CL)	2-89
Copy (CO)	2-90
CO Command Source and Destination Parameter Considerations	2-92
CO Command Options	2-93
CO Command Examples	2-96
Default (DE)	2-99
Directory List (DL)	2-100
Echo Command (EC)	2-102
Exit (EX)	2-102
Group Copy Commands (GR, EG, and AG)	2-102
List Comment, List Header Files (LC, LH)	2-103
List Device (LL)	2-103
Scratch Area Definition (SC)	2-104
Title (TI)	2-104
Transfer (TR)	2-104
Tape Handling	2-105
Destination Disk Handling	2-106
Performance Considerations	2-107
Loading FC	2-108
Using Globals in Transfer Files	2-108
Error Handling in Transfer Files	2-110
FC Error Messages	2-112
Errors Requiring Operator Action/Response	2-112
Information Messages and Warnings	2-113
Disk Data I/O Errors	2-116
Non-Fatal Tape Read Errors	2-116
Errors Affecting a Single File	2-117
Loss of Unidentified Files on Copy from Tape	2-120
Disk-to-Tape Copy Verify Errors	2-121
Errors that Cause Rejection of Current Source Tape Volume	2-121
Command Syntax, Parameter Errors (Command is Skipped)	2-122
Command Out-of-Sequence Errors (Command is Skipped)	2-124

Other Errors that Terminate Current Command	2-125
Other Errors	2-127
Errors that Terminate FC	2-127
FORMA	2-129
FORMA Definitions	2-129
Calling FORMA	2-131
FORMA Commands	2-131
Help (??)	2-131
Initialize Media (IM)	2-131
Spare Sector (SS)	2-132
SS Command Options	2-133
Converting from FORMF to FORMA	2-133
Installing a New Disk	2-133
FORMA Error Handling	2-134
FORMA Error Messages	2-134
User Messages	2-134
User Errors	2-134
Disk Errors	2-135
Disk Error Decoding	2-136
FORMC	2-137
Calling FORMC	2-137
Command Execution	2-138
Device Driver Status	2-138
Break Detection	2-138
FORMC Commands	2-139
Abort, End, and Exit (AB) (EN) (EX) (/E)	2-139
Format (FO)	2-140
Help (?)	2-140
Spare (SP)	2-140
Verify (VE)	2-141
The FORMC Formatting Operation	2-141
Entering the LU	2-141
Disk Formatting	2-141
Tape Formatting	2-144
The FORMC Sparing Operation	2-146
Entering the LU	2-146
Specifying the Track	2-146
Sparing Defective Disk Blocks	2-146
Data Loss During Sparing	2-147
The FORMC Verify Operation	2-148
Entering the LU	2-148
Readying the Tape	2-148
Verifying the Entire Disk LU or Tape	2-148
Specifying the Start Track or Block	2-149
Specifying the Number of Tracks or Blocks	2-149
The Verification Process	2-150
Loading FORMC	2-151
FORMC Error Messages	2-151
FORMF	2-152
Calling FORMF	2-152
Command Execution	2-153
Device Driver Status	2-153
Break Detection	2-153

FORMF Commands	2-154
Abort, End, and Exit (AB) (EN) (EX) (/E)	2-154
Format (FO)	2-155
Help (?)	2-155
Verify (VE)	2-155
The FORMF Formatting Operation	2-156
Entering the LU	2-156
Specifying an Interleave Factor	2-156
The Formatting Process	2-157
The FORMF Verify Operation	2-157
Entering the LU	2-157
Verifying the Disk	2-157
FORMF Error Messages	2-158
FORMT	2-160
Calling FORMT	2-160
FORMT Commands	2-161
End (EN)	2-161
Format a Mini-Floppy Disk (FO)	2-161
Help (??)	2-162
Initialize an LU (IN)	2-162
Reformat a Hard Disk (RE)	2-162
Spare a Track (SP)	2-163
Verify an LU (VE)	2-163
Generating a FORMT System	2-163
The FORMT Formatting Operation	2-164
Entering the LU	2-164
Confirming Formatting	2-164
Sector Interleaving	2-165
The Formatting Process	2-165
The FORMT Initializing Operation	2-166
Entering the LU	2-166
Confirming Initializing	2-166
The Initializing Process	2-166
The FORMT Reformatting Operation	2-167
Before You Use the RE Command	2-167
Entering the LU	2-167
The Reformatting Process	2-167
After You Reformat the Disk	2-167
The FORMT Sparing Operation	2-168
Entering the Track Number	2-168
The Sparing Process	2-168
The FORMT Verify Operation	2-169
Entering the LU	2-169
The Verify Process	2-169
FORMT Error Messages	2-170
FST	2-172
Calling FST	2-173
FST Commands	2-175
Command Stack (/)	2-176
Backup (BA)	2-176
Directory File (DF)	2-177
List Directory (DL)	2-177
SHEMA Size (EM)	2-178

Exit (EX)	2-178
Begin Backup/Restore (GO)	2-178
Help (HE)	2-178
List Comment File (LC)	2-179
List Header (LH)	2-179
List Selected Files (LI)	2-179
Select Log Device/File (LL)	2-179
List Non-Selected Files (LN) (Restore Only)	2-180
Specify Tape LU/Archive File (MT)	2-180
Next (NE)	2-181
Position (PO)	2-181
Previous (PR)	2-181
Restore (RE)	2-182
Run (RU)	2-183
Select Comment File (SC) (Backup Only)	2-183
Set Tape Density (SD) (Backup Only)	2-183
Secure (SE)	2-184
Show (SH)	2-184
TAR (TA)	2-184
Title (TI) (Backup Only)	2-185
Transfer to Command File (TR)	2-185
Unselect (UN)	2-186
FST Options	2-187
Append (A) (Backup Only)	2-188
Brief (B)	2-189
Clear (C)	2-189
Duplicate (D) (Restore Only)	2-189
Faulty (F) (Restore Only)	2-189
Inhibit (I)	2-189
Keep (K)	2-190
Lock (L) (Backup Only)	2-190
MinDir (M) (Restore Only)	2-190
Normal (N) (Backup Only; VC+ Only)	2-190
Original (O) (Restore Only)	2-190
Purge (P) (Backup Only)	2-191
Quiet (Q)	2-191
RwndOff (R)	2-191
SrchApp (S) (Restore Only)	2-191
Update (U)	2-191
Verify (V)	2-192
Whole (W) (Backup Only)	2-192
Yes (Y)	2-192
Z (Z)	2-192
File Masking and Renaming	2-193
D, K, N, and S Qualifiers	2-193
Backing Up	2-193
Restoring	2-195
Incremental Backup	2-197
Restoring from Incremental Backups	2-199
Appending Data	2-201
Consecutive Backups	2-201
Multiple Reels	2-201
Tape Loading	2-202

TF Compatibility	2-202
TAR Compatibility	2-203
UNIX Compatibility	2-203
Rescuing Files from an Overwritten Tape	2-204
Disk Directory File	2-205
Shareable EMA	2-205
Replacing Reserved Characters	2-205
Recommended System Usage	2-206
Streaming	2-206
FST Format	2-207
Installing FST	2-209
FST Error Handling	2-209
FST Error Messages and Warnings	2-210
LIF	2-219
Naming Conventions	2-219
Calling LIF	2-220
LIF Commands	2-220
Copy (CO)	2-222
Directory List (DL)	2-223
Exit (EX)	2-224
Help (HE)	2-224
Initialize (IN)	2-224
List (LI)	2-225
Set Logical List Device (LL)	2-225
Mount Cartridge (MC)	2-225
Pack Cartridge (PK)	2-226
Purge (PU)	2-226
Rename (RN)	2-226
Store (ST)	2-227
Severity (SV)	2-228
Transfer Control (TR)	2-228
LIF Error Handling	2-229
TF	2-230
Calling TF	2-231
TF Commands	2-231
Copy (CO)	2-232
CO Command Source and Destination Parameters	2-233
CO Command Options	2-234
Copy Examples without Subdirectories	2-237
Copy Examples with Subdirectories	2-243
Copy Examples Using DS	2-250
Default (DE)	2-253
Directory List (DL)	2-254
Relation of the DL Command to the CO Command	2-256
Exit (EX)	2-257
Group Copy Commands (GR, EG, and AG)	2-257
Help (HE)	2-258
List Header File (LH)	2-258
List Device (LL)	2-258
Title (TI)	2-259
Transfer Command (TR)	2-260
Tape Protection and the K Option	2-260
Time Stamps	2-261

Create Time	2-261
Update Time	2-261
Access Time	2-261
Missing Time Stamps	2-261
Maintaining the System Time	2-261
Incremental Backup	2-262
Backup Bit, B Qualifier, and C Option	2-262
Standard Incremental Backup Procedure	2-263
Multiple Copies of the Same Backup	2-264
Reusing the Backup Bit	2-264
Restoring Incremental Backups	2-264
Restoring Older Versions from Incremental Backup Tapes	2-265
Setting Backup Bits on Restore	2-265
Alternatives to Standard Incremental Backup	2-266
Replacing Duplicate Files	2-266
Automatic Creation of Directories on Restore	2-267
Saving and Restoring File Properties	2-268
System Backup and Restore	2-269
Handling Disk Full Errors	2-270
File Access During Backup and Restore Operations	2-271
Using TF with FMGR Files	2-271
Using TF with FC Tapes	2-272
Multi-Tape Backup and Restore	2-273
UNIX Compatibility	2-273
File Formats on FMP and UNIX	2-273
Copying Files Between FMP and UNIX	2-274
Summary of Copying Recommendations	2-277
Directory File Names	2-278
File Name Letter Case	2-278
Periods Used in File Names	2-278
Special Characters	2-279
File Name Length—Moving Files from FMP to UNIX	2-279
File Name Length—Moving Files from UNIX to FMP	2-279
Linked Files	2-280
Appending to Tapes	2-280
File Types, Security Codes	2-280
Time Stamps	2-280
Root Directory	2-280
Header and Final Checksum Dummy Files	2-280
TF Tape Format	2-281
Installing TF	2-282

List of Illustrations

Figure 2-1	Example of FC Command Summary Listing	2-88
Figure 2-2	Example of Tape Directory List Format	2-101
Figure 2-3	Typical Track Format	2-130
Figure 2-4	Example Directory Tree Structure	2-194
Figure 2-5	FST Format and Header Basics	2-208
Figure 2-6	Example of Directory List Format	2-255
Figure 2-7	TF Tape Format and Header Basics	2-281

Tables

Table 1-1	Disk Backup Utilities	1-2
Table 1-2	File Interchange Utilities	1-4
Table 1-3	RTE-A Formatting Utilities	1-6
Table 2-1	ARSTR Commands Summary	2-3
Table 2-2	RE Command Options Summary	2-6
Table 2-3	ARSTR LU Restore Output	2-13
Table 2-4	ASAVE Commands Summary	2-20
Table 2-5	SA Command Options Summary	2-23
Table 2-6	ASAVE Header Record Format	2-30
Table 2-7	ASAVE Example	2-33
Table 2-8	DRSTR Commands Summary	2-37
Table 2-9	RE Command Options Summary	2-41
Table 2-10	DSAVE Commands Summary	2-53
Table 2-11	LH Command Example	2-55
Table 2-12	SA Commands Option Summary	2-57
Table 2-13	ERTSH Commands Summary	2-68
Table 2-14	FC Commands Summary	2-87
Table 2-15	FC CO Command Options Summary	2-93
Table 2-16	Responses to System Prompts	2-133
Table 2-17	FORMC Commands Summary	2-139
Table 2-18	FORMF Commands Summary	2-154
Table 2-19	FORMT Commands Summary	2-161
Table 2-20	FST Commands Summary	2-175
Table 2-21	FST Command Options Summary	2-188
Table 2-22	Reserved Character Replacements for TAR Archive Files	2-203
Table 2-23	Reserved Character Replacements for FMGR Files	2-206
Table 2-24	LIF Commands Summary	2-221
Table 2-25	TF Commands Summary	2-232
Table 2-26	TF CO Command Options Summary	2-235

Introduction

RTE-A provides disk backup and formatting utilities that perform a variety of functions. This chapter briefly describes the utilities available. The detailed description of each utility is located in Chapter 2 in alphabetical order.

Backup and File Interchange Utilities

The backup and file interchange utilities available are:

File Storage to Tape (FST) is a high performance backup utility that has streaming capability and supports DDS media, magnetic tapes, and CS/80 cartridge tapes. FST also allows the interchange of files between HP 1000 and UNIX*-based systems.

Tape Filer (TF) is used to copy files between disk and tape media, either disk-to-tape or tape-to-disk. TF is similar to FC (File Copy), but it works with all files, whereas FC works only with FMGR files. TF also allows the interchange of files between HP 1000 and UNIX-based systems.

File Copy (FC) can be used to transfer FMGR files between disk cartridges and tape media, either disk-to-disk, disk-to-tape, or tape-to-disk.

Logical Interchange Format HP Systems File Copy (LIF) allows interchange of files between the HP 1000 and other Hewlett-Packard systems. It translates files from the standard RTE FMP format to a standard logical interface format and vice versa. LIF is similar to the command interpreter (CI); however, it is not intended to be a duplicate of file management capabilities.

Save (ASAVE) and Restore (ARSTR) are used to transfer data between disks and magnetic tape or CS/80 cartridge tape drives (CTD). ASAVE allows you to back up large, contiguous areas of a disk to tape, regardless of the disk's content. ARSTR allows you to restore previously saved disks from tape.

Disk-to-Disk Copy (COPYL) is an offline utility that copies the entire contents of one disk LU to a second disk LU. This utility allows you to make copies of disks without using FMP.

Save (DSAVE) and Restore (DRSTR) are used to transfer data between fixed and removable disks. DSAVE saves selected LUs or an entire disk drive unit to removable disk. The removable disk media include 3 1/2, 5 1/4, and 8-inch floppies and disk platters such as the HP 7906H disk. DRSTR restores the saved data from the removable media to fixed disk.

The physical backup utilities (ASAVE/ARSTR, COPYL, and DSAVE/DRSTR) are not compatible with previous RTE physical backup utilities and, therefore, cannot be used for interchange.

* UNIX is a registered trademark of UNIX System Laboratories Inc. in the U.S.A. and other countries.

Backup Utilities

You may use several utilities to back up and verify data from a disk drive to a tape drive (either magnetic tape, DDS media, or CS/80 cartridge tape). RTE-A supports two modes of backup, file backup and physical disk image backup. Table 1-1 shows the utilities used for each mode.

File (Logical) Backup

File backup allows online backup and restoration to and from disks. It allows selective file backup and selective file restore from a larger backup. Selective file backup can be relatively fast, because you need not back up the entire disk LU. However, you may find that total file backup of a full disk is slower than physical image backup due to the file system overhead. In addition, offline restoration is not possible because file backup requires a functioning, disk-based operating system.

FST and TF are used to back up files from disk to tape and restore files from tape to disk (some restrictions apply to FMGR files). These utilities support incremental backups of just those files that were modified, and they allow you to append new backups to old backups. FST is faster than TF and supports streaming on streaming tape drives. FST can also be used to back up and restore files to and from archive files.

FC is used to back up files from disk to tape and restore files from tape to disk.

All of the above utilities support a Verify option that allows you to check that data was properly saved.

LIF is used primarily to interchange files between Hewlett-Packard systems but can also provide file backup.

CI is the RTE-A command interpreter and can be used to copy files on disk. See the *RTE-A User's Manual*, part number 92077-90002, for a complete description of CI.

Table 1-1. Disk Backup Utilities

	File Backup from Disk	Physical Image Backup from Disk
to Cartridge Tape	FST, TF, FC	ASAVE/ARSTR
to Magnetic Tape	FST, TF, FC	ASAVE/ARSTR
to Disk	CI, FST	COPYL
to DDS Media	FST, TF	ASAVE/ARSTR
to Removable Disk	CI, FST	DSAVE/DRSTR

Physical Backup

Physical backup allows offline restoration and, in the case of a total disk backup, is faster than file backup. However, you cannot restore files selectively using physical backup but must restore the entire disk, which wipes out any previous contents. Physical backup utilities are usually used to save system areas of disks online, so that you can restore these areas offline if a disk becomes corrupt. Transfers are done on continuous tracks because there is less disk head movement; therefore, the utilities tend to run faster than file-oriented backup utilities.

ASAVE and ARSTR are used to back up and restore disk LUs to magnetic tape, CS/80 cartridge tape, or DDS media, regardless of file structure. ASAVE and ARSTR allow you to append new disk LU backups to previous backups and can be used in a memory image system for offline restores and saves.

COPYL can be used to copy the data on one disk LU to another disk LU of similar type, regardless of file structure.

DSAVE and DRSTR are used to transfer data between fixed and removable disks.

Compatibility with Other Disk Backup Utilities

The RTE-A physical backup utilities are not compatible with previous RTE physical backup utilities (for example, PSAVE and PRSTR) or file backup utilities (for example, TF, FC, and LIF). FST is the recommended utility for transporting data between RTE-A systems.

Compatibility Among Disks

RTE-A physical backup utilities do not use file structure information. Therefore, the source and destination disks used in save and restore operations must have the same track size (sectors per track), although the total number of tracks in each disk LU need not match.

If you must transfer data between disk LUs with different track sizes, use one of the file backup utilities (FST, TF, FC, or LIF) to transfer the files individually or in groups, rather than on a track-to-track basis.

The physical backup utilities support all disks that are controlled by RTE-A drivers DD*30, DDM30, DD*33, ID*37, DDQ30, and IDQ35.

File Interchange on RTE-A

You can use the FST, TF, FC, and LIF utilities to interchange files between RTE systems (RTE-A, RTE-6/VM), other Hewlett-Packard systems, and UNIX-based systems. These utilities support a Verify option to ensure the validity of data copied. Table 1-2 shows which utilities to use to move data between systems. The table indicates which utilities are recommended for saving and restoring files from/to a system. For example, “FST/TAR” means that FST is recommended for copying files from one system (the source), and TAR is recommended for restoring files to the other system (the destination).

Table 1-2. File Interchange Utilities

From System \ To System	RTE-A or RTE-6/VM (4) CI files	RTE-A or RTE-6/VM (4) FMGR files	UNIX
RTE-A or RTE-6/VM(4) CI Files	FST/FST	FST/FST	FST/TAR (3)
RTE-A or RTE-6/VM(4) FMGR files	FST/FST	FST/FST	(2)(3) FST/TAR
UNIX	(3) TAR/FST	(1)(3) TAR/FST	-
Notes: (1) File names truncated to 6 characters, time stamps lost (2) Restrictions on use with FMGR files (see TF/FST) (3) UNIX compatibility restrictions (see TF/FST) (4) Revision 2540 or later			

Formatting Utilities

The RTE-A formatting utilities format disk volumes or tape, verify the integrity of disks and CS/80 CTDs, and spare bad areas of media. The formatting utilities available are:

CS/80 Disk Formatting (FORMC) is a maintenance program, provided in both online and offline forms, for CS/80 disk and tape drives.

Disk Formatting (FORMF) is the online maintenance program for floppy and standalone 5-1/4 inch mini-Winchester disk drives.

Offline Disk Formatting and Initialization (FORMT) is an offline utility that formats a floppy disk, identifies bad tracks on the disk, initializes disk LUs, spares bad tracks, verifies a disk, or reformats a disk.

Initialize and Spare (FORMA) is an online utility that provides initializing and sparing functions for Micro/1000 systems with internal disks connected via the HP 12022 disk controller.

Error Rate Test (ERTSH) is an online utility that tests the Micro/1000 internal hard disk.

Table 1-3 lists the disks supported by each utility. Check Table 1-3 before you format a disk, to ensure you are selecting the proper utility for your task.

For SCSI devices, refer to the VSCSI utility in the *HP 12016A SCSI Installation and Reference Manual*, part number 12016-90001.

General Formatting Information

A disk cannot be used (normal reads and writes cannot be made to the disk) until it is formatted. Formatting is the simplest form of initialization. (Note that the terms “formatting” and “initializing” are not interchangeable, as there are differences between the two processes.)

When a disk is formatted, the formatting process writes every block on the disk. All information on a disk is written in blocks of 256 bytes (128 words). The block number is usually called a “sector number” in disk literature; however, because an RTE sector is usually 128 bytes (64 words), the term block is used in this discussion.

Each block has three components, the preamble, the user data area, and the postamble. The disk is written as described below.

preamble	user data – 256 Bytes	postamble
----------	-----------------------	-----------

Each block is tagged with an address (preamble) which initially points to itself, but changes if the block is defective or a spare.

If the disk is normal, the preamble address identifies the block. If the disk is defective, the preamble address is either zero or the address of the next block, depending on the type of drive. In either case, the block is made functionally invisible.

The formatting process writes test data in the user data area, then reads it back to make sure the block is normal. The postamble is written with error checking information calculated from the test data. If the test passes, the block is normal. Each block is tested several times.

Hard disks supplied by Hewlett-Packard are formatted before shipment. Because a hard disk should be formatted only at the factory to ensure interchangeability of disk packs, hard disk formatting capability is not included in the RTE-A formatting utilities. Hard disk initialization assumes the disk was formatted at least once.

Table 1-3. RTE-A Formatting Utilities

FORMC Supports	FORMF Supports	FORMT Supports	FORMA Supports
7907 7908 7911 7912 7914 7933 7935 7942 7945 7946 9122C/D 9133H 9144 CS/80 248x hard disk with floppy, with 12122A	9121 S/D 9133A/9134A 9133B/9134B 9895 9130 K (Model 6 built-in disks) 248x Integrated 3 1/2" microfloppy disk with 12022A	7906 H/M 7910 H/M 7920 H/M 7925 H/M 9130 K (Model 6 built-in disks)	248x Integrated hard disk
Runs under these operating systems:			
Online: RTE-6/VM RTE-A Offline: RTE-A	Online: RTE-A	Online: RTE-6/VM *Offline: RTE-A	Online: RTE-A
* FORMT is supplied only in relocatable form. The instructions in the section on generating a FORMT system must be followed in order to create an offline version of FORMT.			

The Preamble

The preamble has three functions:

- A synchronization field provides timing signals to indicate the start of the block.
- Bits indicate the status of the block. These bits are used differently or not at all depending on the actual disk type (hard or soft) and the disk design.

The preamble is more complex for a hard disk than for a flexible disk. The status portion of the preamble includes three bits:

S bit	Indicates that the track is a spare for a defective track. Only hard disks supplied by HP permit sparing.
P bit	Indicates a protected track, one that is “read only” unless the FORMAT switch is activated on the disk drive.
D bit	Indicates that the track is defective and should not be used.

For flexible disks supplied by HP, only the D bit is used in the preamble.

- An address gives position information. The address field in the preamble has three basic components:

- Cylinder number
 - Head number
 - Block number

In most cases, the address associated with the block is the address of the block itself; that is, it points to itself. The exceptions are:

Hard Disk	If the D bit is set (track is defective), the address points to the track used as a spare.
	If the S bit is set (track is a spare), the address points to the defective track for which it is a spare.
Flexible Disk	If the D bit is set (track is defective), the address is not relevant (may be anything).

The Postamble

The postamble provides information that the disk drive needs for error checking during which time it determines the validity of the data in the user data area.

Initializing and Sparing Hard Disks

Disk initialization uses the track map table in RTE-6/VM or the driver parameters in RTE-A to determine the starting cylinder, head number, number of tracks and number of spare tracks for each LU. Tracks in the spare pool are used only as replacements for the defective tracks and are otherwise unused.

Sparing vs. Skipping

The RTE utility programs declare an entire track bad if one bad block is found on the track during formatting. This method of management results in less head movement (hence faster access) when spared blocks (tracks) are encountered.

On a hard disk, once a track is spared subsequent seeks by the disk controller to that track reveal that the track is spared, and the disk drive automatically accesses the spare track instead of the bad one. This process is transparent to your software.

On a flexible disk, although the technique is not called sparing the effect is almost the same. Once the D bit is set for all defective tracks, they become invisible to your software. This causes the drive to skip bad tracks as if they did not exist on the disk.

For example, the outermost track is normally track zero. If the D bit is set for this track, then a command to access track zero causes this track (and any bad tracks that follow it) to be skipped. The first good track encountered becomes logical track zero.

If no bad tracks occur after formatting is complete, it does not matter whether sparing is provided. Sparing and skipping do not really have any effect on your programs. If a bad track develops when you use the disk, the hard disk sparing capability has the advantage that a simple recovery process is available when you run the FORMT utility. This is not possible with the flexible disk.

Interleave Factor Calculation

When you format a flexible disk, RTE formatting utilities let you define the order in which consecutive logical blocks are located on the disk by specifying an interleave factor or fill number. The interleave factor is the number of physical blocks to skip between logical blocks.

With an interleave factor of zero, block $n+1$ immediately follows block n . If the interleave number is 1, one block is skipped between blocks n and $n+1$. An interleave factor of two causes two physical blocks to be skipped between blocks n and $n+1$, and so on. Once a disk is formatted, the interleave is transparent to your software.

You can calculate the interleave factor using two numbers: the time between disk accesses done by the software, and the rotational speed of the disk. For a program that reads sequential blocks from the disk, the calculation is simple. First the processing time for each block is measured, then the read time per block is calculated. The read time per block equals the time of one disk rotation divided by the number of blocks per track.

For example, suppose the software takes T milliseconds to read and process one block, and it picks up sequential blocks. If the time to rotate once around the disk is divided by the number of blocks per track, the time to move over a single block is found. The interleave number is then computed using the following formula:

$$\text{Processing time} \leq [(\text{interleave}-1) * \text{sector read time}]$$

Thus, the next block rotates into the proper position for access just when it is needed.

Consider a disk that has the following rotation speed and number of blocks per track:

$$\begin{aligned} 360 \text{ RPM (rotations/minute)} &= 2.77 \text{ milliseconds/rotation} \\ 30 \text{ blocks/track means} & 0.09 \text{ milliseconds/block} \end{aligned}$$

Suppose the program takes .15 mS to process each block:

$$.15 \text{ mS per block} / .092 \text{ mS per block} = 1.63$$

The next greater whole integer is 2, so in this case the optimal interleave factor is 2.

If the processing time is less than 0.09 mS, an interleave factor of 2 permits processing to finish before the next block becomes available.

If the processing time is between 0.09 mS and 0.18 mS, an interleave factor of 3 is appropriate.

An interleave factor of 1 means the disk must rotate all the way around before the next block becomes available (2.77 mS vs less than 0.09 mS).

An interleave value of 2 is the smallest value recommended with the current floppy disk controller. An interleave value of 1 causes the disk to wait a full revolution per block when a multiple block transfer is requested, because the transfer rate of the controller does not keep pace with the transfer rate of the disk.

Most user programs that access files define a Data Control Block (DCB) of 144 words, of which 128 words (1 block) are used for disk transfers. Therefore, you can compute a minimum interleave factor in terms of the amount of time to complete one EXEC request to write a block and return to the user program.

Alternatively, you can determine the interleave factor experimentally by trying several interleave numbers. For sequential access to the files with a DCB of 144 words, try an approximate number in the range of 4 to 6.

CS/80 Cartridge Tape Drive Definitions

The CS/80 cartridge tape drive (CTD) is a streaming mode, block record device. CS/80 device driver (DD*33) disk-caching routines, which control transfers to and from the CTD, provide efficient handling of the cartridge tape.

Data to be transferred are buffered on the integrated disk drive and transferred in a steady stream to the CTD, allowing continuous movement of the tape. Refer to the *RTE-A Driver Reference Manual*, part number 92077-90011, for details on the disk caching scheme.

Unless otherwise specified, the following definitions apply to CS/80 disks and cartridge tape drives throughout this chapter.

Block	The smallest unit of information, 128 words in length, that can be addressed by a CS/80 disk. It corresponds to a cylinder, head, or sector address.
Tape	The smallest unit of information, 512 words in length, that can be addressed by a CS/80 tape unit.
Volume	A separately addressable portion of the storage media on a given unit, corresponding to a fixed or removable platter of a disk unit or a tape drive.
Track	A subdivision of a disk LU as defined by the EXEC interface.

Disk Backup and Formatting Utilities

This chapter contains a detailed description of each backup and formatting utility listed in alphabetical order.

ARSTR

ARSTR

ARSTR restores one disk LU, a group of disk LUs, or an entire disk unit from magnetic tape or CS/80 CTD.

When you use ARSTR to restore a disk LU, first dismount the disk volume, run ARSTR, and then mount the volume.

Calling ARSTR

Invoke ARSTR with the following runstring:

```
CI> [RU, ]ARSTR[ ,commands ]
```

If you do not enter a command in the runstring, ARSTR prompts you interactively, as follows:

```
ARSTR :
```

Separate command parameters either with commas or blanks. You may combine multiple commands either in the runstring or in response to the prompt, separated by vertical bars. You may also use the underscore (`_`) character to separate parameters in the runstring; ARSTR converts them later to blanks.

Commands are executed in order from left to right. If any command causes a serious error, execution terminates with that command (all the preceding commands are executed).

ARSTR Commands

ARSTR uses commands that allow you to terminate ARSTR, list all the commands and their parameters, list header information, open a list file or list device, restore a file, rewind and set tape offline, specify a tape LU for ARSTR commands to use, and specify the user error handling mode. Table 2-1 summarizes the commands, which are described in the following sections.

Table 2-1. ARSTR Commands Summary

Commands		Description
Information Command		
HElp		Lists all ARSTR commands and parameters
Configuration Commands		
RE store	<i>file #: lu</i>	Restores files from tape, either an entire LU or a unit restore
RW Rewind		Rewinds and sets the tape offline
TA pe	<i>lu</i>	Sets tape LU to the specified LU or displays current LU
UE User Error Handling	ON OFF	Sets user error handling mode ON or OFF or displays current status
Listing Commands		
LH List Header	<i>file #</i>	Lists header information for the specified file number
LL List Device	<i>file_des lu</i>	Opens a list file or list device to which information is to be logged in addition to the terminal
Exit Commands		
AB ort EN d EX it		These commands all terminate ARSTR

Abort, End, and Exit (AB) (EN) (EX)

Purpose: Any one of these commands terminates ARSTR.

Syntax: AB, EN, or EX

Description:

Upon termination, all disk LUs are released and any list file or device closed.

ARSTR

Help (HE)

Purpose: Lists all ARSTR commands and parameters.

Syntax: HE, ??, or ?

List Header (LH)

Purpose: Lists header information from the tape, for the specified file number.

Syntax: LH[,*file#*]

file# The file for which the header information is being obtained.

Description:

If the specified file does not exist on the current tape, the header record of the first or last file on the tape is listed. If you do not specify a file number, the header record of the next save file is the default. If positioned at the EOD record, the header record of the last save file is listed.

Note

All file numbers specified are relative to the beginning of the first ASAVE tape. If the save file for a disk LU crosses tapes, it has the same file number on both tapes.

List Device (LL)

Purpose: Opens a list file or list device to which terminal displays and error messages are sent in addition to the log device (the terminal).

Syntax: LL[,*file descriptor* | *lu*]

file descriptor The list file.

lu The list device.

Description:

A defined list file is created if it does not already exist, and a list device is opened. Any previously opened list file or device is closed first. If you do not specify a file descriptor/LU, any currently open list file or device is closed.

A duplicate file error occurs if you specify a file and a duplicate file already exists. In this case, the command terminates with an error if user error handling (described below) is ON; otherwise, the command is skipped.

If you specify a device that is down and user error handling is ON, you are prompted to correct the problem and bring up the device, or enter breakmode. If user error handling is OFF, the command is skipped.

If the LU goes down when you use a list device, and user error handling is ON, you are prompted to correct the problem or enter breakmode. If user error handling is OFF, you must enter breakmode. When a break occurs, the list file or device is closed and the command continues to execute.

Restore (RE)

The RE command has two formats, one for LU restores and one for unit restores. Each format is described separately below.

File-by-file (LU) Restore:

Purpose: Specifies an LU restore from tape.

Syntax: RE ,*file# : lu* [, *file# : lu*] [, . . .] [, *options*]

<i>file# : lu</i>	The file number on tape and the disk LU to which it is to be restored. There are two variations of this parameter:
<i>disk lu(: lu)</i>	The file number is defaulted to 1 for the first disk specified, 2 for the second, and so forth.
<i>file number</i>	The destination disk LU (<i>file#</i>) is determined by examining the save definition record of the save file for the disk LU that was saved. If you specify a file number for any of the <i>file# : lu</i> parameters, then all the <i>file# : lu</i> parameters must specify a file number.
<i>options</i>	The options are described in the next section.

Note All file numbers specified are relative to the beginning of the first ASAVE tape. If the save file for a disk LU crosses tapes, it has the same file number on both tapes.

ARSTR

Unit Restore:

Purpose: To specify a unit restore of the disk unit.

Syntax: RE [, *file#* [, *options*]

file# The file number on tape and the destination disk LU.

options The options are described in the next section.

Description:

The restore operation is explained in detail in the section “The Restore Operation” later in this chapter.

RE Command Options

The RE command options parameter allows you to specify one or a combination of the options shown below. Options may be specified in any order. Table 2-2 summarizes the RE command options.

Table 2-2. RE Command Options Summary

Options	Description
Data Records Checksums Not Checked	Checksums are computed for the data records, not checked
UNit Restore	Specifies a unit restore beginning at the file number you enter, or file 1 if you do not enter a number
VERify Restored Disks	Specifies verification of the restored disks

Data Records Checksums Not Checked (NC)

The checksum included in each record written to tape is a one-word complement of the sum of all the words in the data portion of the record.

NC specifies that the data record checksums are not to be checked. Instead, the default is to verify the checksums for the data records that were written with checksums. You may turn off checksum verification for increased performance when the tape device is an HP 7978A reading a 6250-bpi density tape.

Unit Restore (UN)

The UN option specifies a unit restore beginning at the file number you enter, or file 1 if you do not enter a file number. If the file number you enter specifies a save file, it does not have to be the first save file of a unit save operation. The restore operation only restores the specified save file and any subsequent save files if they are from the same disk unit and were saved in the same operation.

Verify Restored Disks (VE)

VE specifies a verification of the restored disks (verification is described below).

Rewind (RW)

Purpose: Rewinds and sets the tape offline. A CTD is also unloaded.

Syntax: RW

Tape LU (TA)

Purpose: Sets the tape LU to the specified LU, or if no LU is specified, displays the current tape LU.

Syntax: TA[,*lu*]

lu The tape logical unit selected.

Description:

If you do not specify an LU, the current tape LU is displayed. Magnetic tape and CTD are the only supported tape devices. The specified tape LU is locked and remains locked until ARSTR terminates or another tape LU is defined.

The tape LU is used by all commands that access the tape. If you do not specify a tape LU, but attempt an operation that requires tape access, an error occurs.

ARSTR

User Error Handling (UE)

Purpose: Sets the user error handling mode or displays the current mode if no parameter is entered.

Syntax: UE[,ON|OFF]

ON You are prompted for corrective action if a recoverable error occurs. If you correct the error, command execution continues. The initial setting is ON.

OFF The default course of action is taken, as described in the discussion of each command.

ARSTR Command Examples

Example 1: Set the list device to LU 6, tape LU to 25, and list the header from file number 2.

```
ll,6|ta,25|lh,2
```

Example 2: Set the list device to LU 12, tape LU to 23, and restore files 1 to 3 to disk LUs 16, 17, and 19 respectively. Verify the restore operation.

```
ll,12|ta,23|re,:16,:17,:19,ve
```

Example 3: Close the previous list device, set the tape LU to 24, and restore files 1 to 3 to disk LUs 16, 17, and 19 respectively. Verify the restore operation.

```
ll|ta,24|re,2:17,1:16,3:19,ve
```

Example 4: Set the tape LU to 25, and restore files 2 and 3 to the disk LUs specified in the save definition records of those files. Verify the restore operation.

```
ta,25|re,2,3,ve
```

Example 5: Set the tape LU to 23, and perform a unit restore starting at file 1. Verify the restore operation.

```
ta,23|re,un,ve
```

The Restore Operation

All disk LUs being restored to are first checked to ensure that they are legal. Checking may occur when you enter the RE command or when the appropriate save file is encountered during the restore, as explained below.

The following sections describe the various elements of a restore operation.

Checking the Disks

Disk LUs entered in the RE command are checked when you enter the command. If any disk LU is illegal, the restore operation terminates at once, and the following error message appears:

```
Error - LU x not a supported disk.
```

Disk LUs that are identified from the data in the save files are checked as the save files are encountered during the restore operation. If any disk LU is illegal for the current system, the following warning message is logged:

```
Warning - LU x not a disk in current system.
```

If user error handling is ON, you are prompted to enter breakmode, skip this restore, or enter the LU of the disk to restore. If user error handling is OFF, the restore is skipped.

Sectors per Track Must Match on Disks

ARSTR only restores a disk save file on tape to a similar disk. For the disk to be similar, the sectors per track of the saved disk must be the same as those of the disk to be restored to. Otherwise, the following error occurs, and the restore of that disk is skipped:

```
Error - Sectors per track mismatch.
```

Total Number of Tracks May Differ

The total number of tracks of the saved disk (source) may differ from the total number of tracks on the target disk (destination). If the number of tracks of the destination disk differs from the source disk, ARSTR issues a warning message. If user error handling is ON, ARSTR prompts you to confirm the restore. For example:

```
Warning - Saved disk has less tracks than target disk LU x
Saved disk tracks = x
Target disk tracks = x
Restore this disk ?
```

If user error handling is OFF (the default), the following message appears, and the restore operation continues:

```
Continuing to restore.
```

ARSTR

If the source disk has fewer tracks than the destination disk, tracks beyond where the restore completes remain unaltered. The destination disk LU can be used as a holding area but cannot be mounted because the bitmap of the source disk LU and the destination LU may not be the same.

If the destination disk has fewer tracks than the source disk, the restored source tracks will be limited to the number of tracks on the destination disk.

Disk Locking

Each disk LU being restored is locked when the restore operation begins, to ensure the integrity of the restored data. It remains locked until the restore completes (if the VE option was not specified), or until verification completes (if VE was specified). If the disk cannot be locked (for example, it is currently locked to another program) and user error handling is ON, you are prompted with the following message:

```
Error - LU x locked to (program name) Correct problem and type
GO to continue, SK to skip or BR to break.
```

If user error handling is OFF, the restore of that disk is skipped.

Order of Disk Restoration

Disks are restored in the same order as they were saved to tape, regardless of the order in which they may be specified in the RE command.

Unit Restore

For a unit restore, the tape is positioned to the file number specified in the RE command. ARSTR restores the save file and any succeeding save files that are on the same disk unit and were saved in the same save operation. Information contained in the save definition record determines if these conditions are met.

Save Definition Records

Information for the restore of a disk is obtained from the system and the save definition record in the ASAVE file. Data in the save definition record is used to control the restore.

Each tape save file being restored is checked to ensure that it is a legal ASAVE file. If it is not, the restore operation terminates with the following message:

```
Error - Non-ASAVE tape.
```

When the actual restore begins, the header record of the save file is logged. If a checksum error occurs on the header record, the following warning displays, and the restore continues:

```
Warning - Header record checksum.
```

If the checksums of both the save definition records fail to verify, the restore of that file is skipped. When this occurs the following error is logged:

```
Error - Save definition records checksums.
```

Restoring the Data

If the save file being positioned to (LU or unit restore) is not on the currently mounted tape, if user error handling is ON, you are prompted with this message:

```
Error - First file on tape is x searching for file y. Correct
problem and type GO to continue, SK to skip or BR to break.
```

If user error handling is OFF, the restore of the disk LU is skipped.

For each save file, disk data from the tape is restored to the disk LU until the end of the disk LU, or the end of the disk data in the save file is reached. If checksum errors occur on data records during the restore, the read of that record is retried up to five times before the following error message is logged:

```
Error - Data record checksum.
Data targeted for disks blocks x to y.
Continuing.
```

When the restore of a disk LU completes, the following message is logged:

```
Restore of file x to disk LU y complete.
```

If the end-of-tape condition is detected before the end of the save file is detected, the restores from the current tape are verified (if VE was specified), and you are prompted with the following message:

```
Mount next tape and type GO to continue or BR to break.
```

If a tape LU or disk LU goes down during the restore, the following message displays:

```
LU x is down.
Correct problem and UP device, or use BR to break.
```

The system may also issue messages about the device that is down.

Verification

Restore verification is similar to save verification. When you select the VE option for a tape-to-disk restore, the restore operation becomes a two pass process. During the first pass (restore), data is read from the tape, checksums (if written with the data) are verified, and the data is written to disk. In the second pass (verify), the data is re-read from tape into one buffer. The corresponding data restored to disk is read into a second buffer, and the data in the two buffers are compared bit-by-bit.

ARSTR

If a mismatch between the tape and disk occurs, the following warning appears, and the verify of that restore continues to the end of the disk or until three verify errors occur:

```
Warning - Possible data loss in disk block range x to y.
```

If three verify errors occur, the verification process aborts.

Break Detection

If ARSTR detects a break, it terminates all operations and unlocks all LUs except tape LUs. If a break occurs during a restore or during list header command execution, the tape is rewound.

Loading ARSTR

Load ARSTR online, using the LINK loader and the LINK command file #ARSTR.

Offline System

Because ASAVE and ARSTR are not segmented, they may be loaded into a memory-based system for offline use. The memory-based, offline system serves three purposes: installing systems, restoring corrupt system disk LUs, and saving disk LUs from crashed disks.

In general, after you establish a new system, you should use the same system file to create an offline restore system. This should include a startup program (for example, COMND * or CI) that enables the scheduling of programs; the restore (ARSTR) utility, and optionally, the save (ASAVE) utility; and the appropriate format (VSCSI, FORMC, FORMT, FORMF) utilities. Using the same system file as the current system eliminates any need to reconfigure the offline system.

The BUILD utility is used to build the actual memory-based offline system. The following example illustrates a BUILD command file:

```
YES ,           (Automatic partitioning)
256 ,          (Memory size in pages)
RP , ASAVE , , (Relink and create ID segment for ASAVE)
RP , ARSTR , , (Relink and create ID segment for ARSTR)
RP , FORMC , , (Relink and create ID segment for FORMC)
RP , D.RTR , , (Relink and create ID segment for D.RTR)
RP , COMND , , (Relink and create ID segment for COMND)
ST , , , , , (Make COMND the startup program)
/E           (Complete the build process)
```

After you create the system, place it onto bootable media. For detailed instructions on assembling a memory-based, offline system, refer to the *RTE-A System Generation and Installation Manual*, part number 92077-90034.

* COMND is a compact version of File Manager (FMGR) that supports only a subset of FMGR commands. See the FMGR documentation in the *RTE-A User's Manual* for more information.

ARSTR Example

Table 2-3 shows the output of a LU restore. Here the list file 'RSTR is created, the tape LU is defined as LU 8, and a restore with verify of save files 1 and 2 is performed. You must dismount the disk LUs to be restored before, and mount them after running ARSTR. When the restore completes, the tape is rewound and set offline so it cannot be overwritten.

Table 2-3. ARSTR LU Restore Output

```

CI> DC 12
CI> DC 57
CI> ARSTR

ARSTR - REV.6000 <920903.1620>
Use ? for help.
ARSTR : ll,'rstr
Opened list file.
ARSTR : ta,8
Tape LU = 8
ARSTR : re,1,2,ve

Created Using: ASAVE 92077-16586 REV.6000 <920903.1410>
Tape Set Date: Tue Dec 1, 1992 3:20:41 pm
Tape #: 1      File #: 1      Options: VE      Disk LU: 12
Tracks: 50    Sec/Trk: 96    Section: 1    Trk: 0      Sec: 0
Created: Tue Dec 1, 1992 3:20:41 pm      User: SYSTEM
Title: Sample LU save

Restore of file 1 to disk LU 12 complete.

Created Using: ASAVE 92077-16586 REV.6000 <920903.1410>
Tape Set Date: Tue Dec 1, 1992 3:20:41 pm
Tape #: 1      File #: 2      Options: VE      Disk LU: 57
Tracks: 48    Sec/Trk: 96    Section: 1    Trk: 0      Sec: 0
Created: Tue Dec 1, 1992 3:21:04 pm      User: SYSTEM
Title: Sample LU save

Restore of file 2 to disk LU 57 complete.
Verifying restores from tape.
Verifying restore of file 1 to LU 12
Verifying restore of file 2 to LU 57
ARSTR : rw
ARSTR : ex
Closed list file.

CI> MC 12
CI> MC 57

```

ARSTR

ASAVE and ARSTR Error Messages and Warnings

The sections that follow list 1) all the messages that ASAVE generates, 2) error messages that are shared by ASAVE and ARSTR, 3) shared warning messages, and 4) the error messages that ARSTR generates.

ASAVE Messages

Error – Can only append to last tape.

You selected AP when saving to a tape that was not the last tape in a tape set. Mount the last tape, and retry the operation.

Error – Corrupt tape, no EOD.

A corrupt tape resulted when an ASAVE break was unable to clean up. Perhaps the tape LU was down or not ready, or you aborted ASAVE. Only the last file is corrupt. You may not append to this tape.

Error – Data record verify failed.

A bit-by-bit compare of the tape data record to the disk data detected an error (this also verifies the checksum).

Error – EOD checksum bad, unable to append.

A checksum error was detected in the EOD record while positioning to it for an append save. Since information in the EOD record is used in the save operation, and the correctness of that data cannot be guaranteed, the append save is not allowed.

Error – EOF verify failed.

An EOF failed to verify during the verify pass of a save operation.

Error – Header record verify failed.

A checksum error was detected during verification of the header record of a save file.

Error – Illegal unit save.

The specification for the unit save contained more than one disk LU.

Error – No disk to save.

The SA command did not specify any disk LUs.

Error – Save definition record verify failed.

A checksum error was detected during verification of one of the save definition records of a save file.

Error – Tape LU *xx* write protected.

The SA command required write access to a write-protected tape LU.

ASAVE and ARSTR Shared Error Messages

Error – File number.

You specified an illegal file number in either the RE or LH command, or you may have specified the file number more than once, which is not allowed. In the RE command, if one file number is specified for a *file# : lu* parameter, all *file# : lu* parameters must specify a file number.

Error – First file on tape is *xx*

or

Error – Last file on tape is *xx*

The specified file cannot be found on the current tape. This error may be followed by additional information about the error. The corrective action may be to mount the correct tape.

Error – *<fmp error message>*

An FMP error occurred while accessing the list file or device.

Error – Illegal tape LU.

You used the TA command to enter a tape LU that is illegal in the current system.

Error – IO*xx* on LU *xx*

A device I/O error was detected.

Error – List LU = *x*

You tried to define the tape LU as the currently defined list device.

Error – LU *xx* is a mirrored volume

The LU you entered is a mirrored volume LU. You must specify the physical paired disk LU associated with the mirrored volume LU. (Refer to the *Data Pair/1000 Reference Manual*, part number 92050-90001, for more information.)

Error – LU *xx* locked to *<program name>*

An attempt to lock the LU failed because the LU is locked to another program. This error can occur for both tape and disk LUs.

Error – LU *xx* not a supported disk.

The LU you entered is not a legally supported disk, or it is not a disk LU in the current system.

Error – Non-ASAVE tape.

The tape being accessed is not an ASAVE format tape.

ARSTR

Error – Not enough memory (size program).

The utility does not have enough addressable memory for its buffers. You must size the program.

Error – Save definition records checksums.

A checksum error was detected while reading the save definition records of a save file (both save definition records have checksum errors).

Error – Tape retries exhausted.

or

Error – Disk retries exhausted.

The available retries of a tape or disk access are used up. The tape or disk media may be deteriorating.

Error – Tape LU undefined.

A command requiring tape access was entered, but no tape LU was defined. Specify the tape LU with the TA command, and reenter the command.

Error – Tape LU = *x*

You entered the currently defined tape LU as the list device.

Error – Unknown command.

An unknown command was entered.

Error – Unknown parameter.

An unknown parameter was encountered in a command.

Error – Uninitialized tape.

An uninitialized CTD tape was detected. Initialize the tape with the FORMC utility.

ASAVE and ARSTR Shared Warning Messages

Warning – CTD may be loading.

A tape-not-ready error may occur if you attempt to access the CTD tape while the tape is loading. Wait until the tape completes loading, then continue.

Warning – Data record checksum. Data targeted for disk blocks *xxx* to *xxx* Continuing.

A checksum error occurred on a data record being restored. This happens only after the read from the tape is retried five times. If you selected the VE option, this disk block range may be smaller during the verify pass. In any case, the restore of that LU continues.

Warning – Header record checksum.

There was a checksum error reading the header record. Because no information in the header is used by the utility, this is not a serious error.

Warning – Last file corrupt.

The break during a save operation could not clean up, because the tape LU could not be accessed (it was down, or not ready).

Warning – LU *xx* is a mirrored volume.

An LU that was determined from the save file during a restore operation is a mirrored volume LU in the current system. (Refer to the *Data Pair/1000 Reference Manual* for more information.)

Warning – LU *xx* not a disk in current system.

The LU named is not legal in the current system.

Warning – Possible data loss in disk block range *xxx* to *xxx*

The restore verify pass detected the possibility of data loss in the specified disk block range when it did a bit-by-bit comparison of the tape data and disk data.

Warning – Saved disk has more tracks than target disk LU *xx*

or

Warning – Saved disk has less tracks than target disk LU *xx*

There is a difference between the total number of tracks of the saved disk on tape and the disk you are restoring to. If user error handling is ON, you must confirm that you want the restore to continue.

ARSTR

ARSTR Error Messages

Error – Bad LU.

A bad LU specification in a restore command was detected.

Error – Illegal unit restore.

You made an illegal specification for a unit restore; perhaps you entered more than one file number.

Error – No LU or file number.

You entered a restore command for an LU without specifying the file number, or specified the LU as a parameter to select a file to be restored.

Error – Sectors per track mismatch.

There is a difference in the sectors per track of the disk saved on tape and the disk to which you are trying to restore. The sectors per track for the saved disk on tape and the disk you are restoring to must be the same.

Error – Error msg – Wrong Tape.

This message appears if you try to restore a series of tapes in the wrong order.

ARSTR Warning Messages

Warning – Bad Checksum.

This message may be ignored if your restore continues and the tape verifies okay.

ASAVE

ASAVE saves data from hard disks either to magnetic tape or CS/80 cartridge tape drives. ASAVE allows you to back up large, contiguous areas of disks to tape regardless of the content of the disks (such as files). ASAVE saves one disk LU, a group of disk LUs, or an entire disk unit to tape. You may restore the saved disk areas later, with ARSTR.

Caution ASAVE and ARSTR must be of the same revision. If they are not, you will receive the following error message:

```
Warning - Header record checksums
Error - Save definition records checksums
Restore skipped.
```

The exception to this is that at RTE-A Revisions 5010 and 5020 the compatible revisions of ASAVE and ARSTR are as follows:

At RTE-A Revision:	5010	5020
ASAVE	5010	5010
ARSTR	5000	5020

Calling ASAVE

Invoke ASAVE with the following runstring:

```
CI> [RU, ]ASAVE[ ,commands ]
```

Alternatively, you may run ASAVE interactively, entering commands at the following prompt:

```
ASAVE :
```

Separate command parameters either with commas or blanks. You may combine multiple commands either in the runstring or in response to the prompt, separated by vertical bars.

Commands are executed in order, from left to right. If a command causes a serious error, execution terminates with that command (all the preceding commands are executed).

ASAVE

ASAVE Commands

ASAVE uses commands that allow you to terminate ASAVE, list all the commands and their parameters, list header information, open a list file or list device, rewind and set tape offline, specify a disk-to-tape save operation, specify a tape LU for ASAVE commands to use, set the title to use in save operations, and specify the user error handling mode. Table 2-4 summarizes the commands, which are described in the following sections.

Table 2-4. ASAVE Commands Summary

Commands	Description
Information Commands	
HElp	Lists all ASAVE commands and parameters
Configuration Commands	
REwind	Rewinds and sets tape offline
SAve <i>lu</i>	Specifies a disk-to-tape save operation
TApe <i>lu</i>	Specifies the tape LU to use
TItle <i>text</i>	Sets the title for subsequent save operations
UE User Error Handling ON OFF	Selects the user error handling mode or displays current mode
Listing Commands	
LH List Header <i>file #</i>	Lists header information from the tape
LL List Device <i>file_desc lu</i>	Opens list file or device to which information is to be logged in addition to the terminal
Exit Commands	
ABort ENd EXit	These three commands terminate ASAVE

Abort, End, and Exit (AB) (EN) (EX)

Purpose: Any one of these commands terminates ASAVE.

Syntax: AB, EN, or EX

Description:

When the program terminates, the list file or device specified is closed, and all locked LUs are released.

Help (HE)

Purpose: Displays a list of all ASAVE commands and the parameters of each command.

Syntax: HE, ??, or ?

List Header (LH)

Purpose: Lists the header information for the specified file number from the tape.

Syntax: LH[,*file#*]

file# The file for which the header information is being obtained.

Description:

If the specified file does not exist on the current tape, the header record of the first or last file on the tape is listed. If you do not specify a file number, the header record of the next save file is the default. If positioned at the end-of-data (EOD) record, the header record of the last save file is listed.

Note

If there is more than one ASAVE tape, all file numbers specified are relative to the beginning of the first tape. If the save file for a disk LU is on more than one tape, it has the same file number on both tapes.

ASAVE

List Device (LL)

Purpose: Defines a list file or list device to which error messages and terminal displays are to be sent in addition to the terminal.

Syntax: LL[, *file descriptor* | *lu*]
file descriptor The file to be opened.
lu The list device.

Description:

If the list file you define does not already exist, it is created, and a list device is opened. (A previously opened list file or device is closed before the new list file is created or the list device is opened.) If you do not specify a file descriptor/LU, the currently defined list file or device is closed.

A duplicate file error occurs if you specify a file that already exists. In this case, the command terminates with an error if user error handling (described below) is ON; otherwise, the command is skipped.

If you specify a device that is down and user error handling is ON, you are prompted to correct the problem and bring up the device or enter breakmode. If user error handling is OFF, the command is skipped.

If the LU goes down when you use a list device, and user error handling is ON, you are prompted to correct the problem and bring up the device, or enter breakmode. If user error handling is OFF, you must enter breakmode. When a break occurs, the list file or device is closed, and the command continues to execute.

Rewind (RW)

Purpose: Rewinds and sets the tape offline. A CTD is also unloaded.

Syntax: RW

Save (SA)

Purpose: Specifies a disk-to-tape save operation.

Syntax: SA, *lu* [, *lu* [, . . .]] [, *options*]
options The options available for use with the SA command are described in the next section.

Description:

The save operation is explained in detail in the section “The Save Operation” which follows the discussion of the remaining ASAVE commands.

SA Command Options

The SA command options parameter allows you to specify one or more of the options shown below, in any order. Table 2-5 summarizes the SA command options.

Table 2-5. SA Command Options Summary

Options	Description
AP pend	Appends save files to current ASAVE tape
Data Records Not Checked	Checksums are computed for the data records, not checked
No Locking	Disk LUs are not locked during save operation
Duplicate	Replaces duplicate files
UN it Save	Specifies a unit save of the disk pointed to by the first LU parameter
Disk-to-Tape VE rify	Performs a disk-to-tape verification of the save operation

Append to Tape (AP)

AP appends the save files to the current ASAVE tape. The tape must be the last tape in an ASAVE tape set.

Data Record Checksums Not Checked (NC)

The checksum included in each record written to tape is a single-word one's complement of the sum of all the words in the data portion of the record.

NC specifies that data record checksums on the tape are not to be checked. Instead, the checksum is computed for the data records. You may turn off the computation for increased performance when the tape device is a 7978A recording with 6250-bpi density.

No Locking (NL)

NL specifies that the disk LUs are not locked during the save operation. The default is to lock the disk LUs during the save to ensure the integrity of the saved disk data.

If you select both NL and VE, errors may occur during the verify pass if any disk data was altered (written) between the save and verify passes while the disk was unlocked.

ASAVE

Unit Save (UN)

UN specifies a unit save of the disk unit pointed to by the first LU parameter. If you do not enter UN, the save is an LU save of all the LUs specified. See also the section “The Save Operation” later in this chapter.

Disk-to-tape Verify (VE)

VE performs a disk-to-tape verification of the save operation. See the description of the verification method in the section called “Verification” later in this chapter.

If you specify both NL and VE, errors may occur during the verify pass if any disk data was altered (written) between the save and verify passes while the disk was unlocked.

Tape LU (TA)

Purpose: Specifies the tape LU to be used by the ASAVE commands.

Syntax: TA[, *lu*]

lu The tape logical unit to be used.

Description:

If you do not specify an LU, the current tape LU is displayed. Magnetic tape and the CS/80 CTD are the only supported tape devices. The specified tape LU is locked and remains locked until ASAVE terminates or you specify another LU. If you do not specify a tape LU, but attempt an operation that requires tape access, an error occurs.

Title (TI)

Purpose: Sets the title to use in subsequent save operations.

Syntax: TI[, *text*]

text The text for the title, up to 40 characters.

Description:

The title you specify is placed in the header record of the save files on tape. Any characters over 40 are truncated. If you do not enter any text, the title consists of blanks.

When you enter TI in the runstring, be aware that CI interprets both blanks and commas as parameter delimiters. To get around this, ASAVE allows you to use the underscore (_) character to replace blanks in titles in the runstring. ASAVE then converts the underscores to blanks.

CI also converts all lowercase characters to uppercase in the runstring before ASAVE retrieves it. For example:

```
CI> asave|ti Unit save of LU 37 without VE
ASAVE : ti
Title = UNIT SAVE OF LU 37 WITHOUT VE
ASAVE :
```

However, when you enter the title interactively, it remains the same as typed. For example:

```
ASAVE: ti Unit save of LU 37 without VE
Title = Unit save of LU 37 without VE
ASAVE :
```

User Error Handling (UE)

Purpose: Sets the user error handling mode or displays the current mode.

Syntax: UE[,ON|OFF]

ON You are prompted for corrective action if a recoverable error (such as a typing error) occurs. The initial setting is ON.

OFF The default action is taken, as described in the following sections.

Description:

If you correct the error, command execution continues. Error conditions and the default actions are explained in the section “The Save Operation”.

You may also use UE to see whether user error handling mode is ON or OFF by entering the command without any parameter. For example:

```
ASAVE: ue (Displays current UE mode)
User Error Handling = ON
ASAVE: ue off (Turns UE OFF)
User Error handling = OFF
ASAVE :
```

ASAVE

ASAVE Command Examples

Example 1: List the help information to LU 6.

```
11,6|he
```

Example 2: Set the list device to LU 6, tape LU to 25, specify title “Example unit save”, and perform a unit save of the disk unit pointed to by LU 16. Lock the disks during the save and verify the save.

```
11,6|ta,25|ti,Example unit save|sa,16,un,ve
```

Example 3: Set the list device to LU 12, tape LU to 24, title to “Example LU save”, and save disk LUs 16,17, and 19. Keep the disks unlocked during the save and verify the save.

```
11,12|ta,24|ti,Example LU save|sa,16,17,19,ve,nl
```

Example 4: Set the list device to LU 6, tape LU to 25, title to “Example unit and LU save.” Do a unit save of the disk unit pointed to by LU 16, then append an LU save of LU 30 to the same tape. Lock the disks during the save, and verify the save operation.

```
11,6|ta,25|ti,Example unit and LU save|sa,16,un,ve|sa,30,ve,ap
```

The Save Operation

When you use the SA command, ASAVE first verifies that all the disk LUs are legal in the current system. If any LU is illegal, the command terminates with the following error logged:

```
Error - LU x not a supported disk.
```

The sections that follow describe the various elements of a save operation.

Unit Save LU Checking

If you select the UN option to specify a unit save operation, ASAVE locates all the disk LUs associated with the disk unit pointed to by the disk LU you enter, then proceeds in the same manner as for an LU save.

Disk Locking

If you do not select the NL option with the SA command, the disk LUs being saved are locked (the default). If you do not select VE, the disk LU is locked just before it is saved and remains locked until the save completes. Selecting VE causes the disk LU to remain locked until after the save is verified.

If the disk cannot be locked (it is currently locked to another program), when user-error handling is ON, you are prompted with the following message:

```
Error - LU x locked to (program name)
Correct problem and type GO to continue, SK to skip or BR to
break.
```

If user error handling is OFF, the SA command is skipped, and the disk LU is not saved.

Tape Positioning

Saving disks to tape may begin in one of two places on the tape, depending upon whether you select the AP option. If you do not use AP, the save begins at the beginning of the tape and overwrites any data previously written there.

When you select AP, the save files are appended to the current ASAVE tape. First, the current data on the tape is checked to make sure it is an ASAVE tape. If it is not, the command terminates with the following error:

```
Error - Non-ASAVE tape.
```

If it is an ASAVE tape, the EOD record is located (see the section on “ASAVE Tape and File Formats” later in this chapter). The EOD record marks the end of save files on the current tape and specifies whether the save operation continued onto another tape. If it did cross tapes, you cannot use AP in this save operation, and the save terminates with the following error logged:

```
Error - Can only append to last tape.
```

If the EOD record shows this is the end of the tape set, the save begins after the last save file on the tape.

If you are using a CTD, an end-of-file (EOF) mark is written in block 0 of the tape, and the save operation begins at block 1.

Save File Records

ASAVE save files hold three types of records: An ASCII header record at the beginning of the save file that contains user information about the save; two save definition records, and one or more disk data records that contain images of portions of the disk being saved. All records written to tape contain checksums. A more detailed explanation of these records and the ASAVE tape format is provided in the section called “ASAVE Tape and File Formats” later in this chapter.

Saving the Disks to Tape

When the save operation begins, the disk is locked unless you specified NL. The header record, save definition records, and disk data records are written to the tape. The save proceeds either until the end of the disk LU is reached or there is not enough usable tape left to write the next disk data record.

ASAVE

If too little tape remains, the appropriate EOD record is written on the tape, followed by two EOF marks. If you selected the VE option, the portion of the save operation on the tape is verified, the tape is rewound, and the following message is displayed:

```
Mount next tape and type GO to continue or BR to break.
```

The last save file written to the tape is also followed by an EOD record, indicating the end of the save, followed by two EOF marks. This constitutes the end of a save operation on the tape.

If the tape or the disk LU goes down during the save operation, the following message appears:

```
LU x is down.  
Correct problem and UP device or use BR to break.
```

You may also see system error messages about the device that is down.

Verification

When you select the VE option for a disk-to-tape save, the save becomes a two pass process. During the first pass (save), the header and save definition records are written to tape with checksums, the data blocks are read from disk, an NC option-setting-dependent-checksum is calculated for the data blocks, and the data blocks and checksums are written to the tape.

In the second pass (verify), the tape is repositioned to where the save began on the tape, and the following messages are logged:

```
Verifying tape.  
Verifying save of LU x to file x      (For each save file)
```

For each save file, checksums written with the header and save definition records are verified. If a checksum error occurs during verification, the appropriate error message is logged. For example:

```
Error - Header record verify failed.  
or  
Error - Save definition record verify failed.
```

Data records are verified by reading the data records from tape into one buffer and the corresponding data on the disk into a second buffer; then the two buffers are compared bit-by-bit. If a data record fails verification, the following error is logged:

```
Error - Data record verify failed.
```

File marks and the EOD record are also verified, with the following error logged:

```
Error - EOF verify failed.  
or  
Error - EOD verify failed.
```

If a verify error occurs on a save file, the appropriate error is logged and the rest of the file is skipped. The verify continues with the next file.

Break Detection

When a break is detected, ASAVE terminates any operation and unlocks all locked LUs except tape LUs.

If a break is detected during a save operation, an attempt is made to “clean up” by placing an EOD record after the last save file that successfully completed. If this cannot be done (for example, the tape LU is down or you aborted ASAVE), the result is a corrupt ASAVE tape. If the tape LU is down, the following warning message is logged:

```
Warning - Last file corrupt.
```

The tape is corrupt in that it does not contain an EOD record, and the last file is incomplete. You may not append to the tape, and although previous files may be restored with no problems, if you try to restore the corrupted file, unpredictable results may occur.


Loading ASAVE

Load ASAVE online, using the LINK loader and the LINK command file #ASAVE.

ASAVE Tape and File Formats


ASAVE places an EOF mark in the first block (block 0) of any CTD tape that contains ASAVE data. This provides protection from possibly overwriting a CS/80 disk if a front panel pushbutton restore is attempted with an ASAVE CTD tape. (In that case, the EOF in the first block is read, and since the EOF signifies the end of the file, the restore operation terminates immediately, and no data is overwritten on the disk.)

The ASAVE CTD tape format for a single LU save is shown below. The ASAVE magnetic tape format is identical, except it does not contain the EOF at the beginning of the tape.



A unit or multiple LU save is a series of single LU saves followed by the EOD record and EOFs. A unit save starts with the lowest numbered LU on the unit and proceeds through the highest numbered LU. In a multiple LU save, the LUs are saved in the order in which they were specified in the SA command.

The ASAVE CTD tape format for unit and multiple LU saves is shown below. The ASAVE magnetic tape format for unit and multiple LU saves is identical, except it does not contain the EOF at the beginning of the tape.



ASAVE

ASAVE Tape Records

The ASAVE file has the following record format:

Header	Save Def 1	Save Def 2	Data	...	Data	EOF
--------	------------	------------	------	-----	------	-----

Header Records

The ASCII header record is the first record in an ASAVE file. Table 2-6 shows the format of the header record, followed by descriptions of the various fields in the header.

Table 2-6. ASAVE Header Record Format

<pre>Created Using: ASAVE 92077-16586 REV.6000 <920903.1410> Tape Set Date: xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx Tape #: xxx File #: xxx Options: xx xx xx xx disk LU: xxx Tracks: xxxxx Sec/Trk: xxx Section: xx Trk: xxxxx Sec: xxx Created: xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx User: xxxxxxxxxxxxxxxxxxxx Title: xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx</pre>	
Created Using	<i>The version of ASAVE used to save the disk.</i>
Tape Set Date	<i>The time stamp when the first file of the tape set was created; used to identify the various volumes of a tape set.</i>
Tape #	<i>Number of which tape this section of the save file is located.</i>
File #	<i>Save file number.</i>
Options	<i>Options selected for the disk LU save.</i>
disk LU	<i>Disk LU saved.</i>
Tracks	<i>Total number of tracks for the disk LU.</i>
Sec/Trk	<i>Sectors per track for the disk LU (64 word sectors).</i>
Section	<i>Save file section number, used when a save file crosses tapes.</i>
Trk	<i>Disk track where disk data in this section of the save file starts.</i>
Sec	<i>Sector in track described above.</i>
Created	<i>Time stamp when disk was saved.</i>
User	<i>User account name where ASAVE was scheduled.</i>
Title	<i>Title specified for the save operation.</i>

Save Definition Record

Save definition records are located in the second and third records of the save file. The second save definition record is a copy of the first. Their format is shown below.

Word	Content
1	Unused
2	Checksum
3	Length of save definition record
4	Unused (reserved)
5	Tape number
6	File number (relative to beginning of tape set)
7	File number (relative to beginning of current save operation)
8	Number of saves remaining (for the save operation)
9	Type of save (1=LU, 2=UNIT)
10	Operating system (.OPSY)
11	Size of data records (less 2 word checksum)
12–13	Time stamp
14–15	Block # where previous save file starts (CTD only)
16–17	Tape set ID
18–31	Unused (reserved)
32	Disk LU
33	Disk type
34	Save section number
35	Starting track
36	Starting sector
37	Disk address
38	Disk unit
39	Disk volume
40	Cylinder
41	Head
42	#Surfaces
43	#Tracks
44	#Spares
45	#64 word sectors/track
46	Reserved
47–64	Unused (reserved)

Data Record

All the save file data records are of a uniform length, except the last data record may be shorter. The first two words of the data record make up a checksum. The third and fourth words are a disk block address where the data starts. Word five is the file number. Words 6 through 10 are unused, and disk data begins in word 11.

ASAVE

End-of-Data Record

The EOD record marks the end of the save files on the current tape and contains information that specifies whether the save continues on another tape. Its format is shown below:

Word	Content
1	Unused
2	Checksum
3	EOD record length
4	Type of EOD (described below)
5	File # of last save file (relative to the save operation)
6	File # of last save file (relative to the beginning of tape set)
7	Tape #
8–9	Block # where last save file starts (CTD only)
10–14	'ASAVE EOD'
15–16	'END' or 'CONT'
17–18	Time stamp
19–20	Tape set ID
21–32	Unused (reserved)

The “type of EOD” (word 4) may have one of the three values below:

- 1 = end of save operation
- 2 = saving the current disk LU continues on the next tape
- 3 = save continues on next tape with a new disk LU

ASAVE Example

Table 2-7 shows the output (to the log device) of an LU save. In this example, a list file save is created, the tape LU is defined as LU 8, the title is set to “Sample LU save,” and a save with verification of disk LUs 12 and 57 is done. After the save operation completes, the tape is rewound and set offline so other users cannot overwrite it.

Table 2-7. ASAVE Example

```

ASAVE - REV.6000 <920903.1410>
Use ? for help.
ASAVE : ll,'save
Opened list file
ASAVE : ta,8
Tape LU = 8
ASAVE : ti, Sample LU save
Title = Sample LU save
ASAVE : sa,12,57,ve
Disk LU(s) to be saved:
12
57

Created Using: ASAVE 92077-16586 REV.6000 <920903.1410>
Tape Set Date: Tue Dec 1, 1992 3:20:41 pm
Tape #: 1      File #: 1      Options: VE      Disk LU: 12
Tracks: 50    Sec/Trk: 96    Section: 1      Trk: 0      Sec: 0
Created: Tue Dec 1, 1992 3:20:41 pm      User: SYSTEM
Title: Sample LU save

Created Using: ASAVE 92072-16586 REV.6000 <920903.1410>
Tape Set Date: Tue Dec 1, 1992 3:20:41 pm
Tape #: 1      File #: 2      Options: VE      Disk LU: 57
Tracks: 48    Sec/Trk: 96    Section: 1      Trk: 0      Sec: 0
Created: Tue Dec 1, 1992 3:21:04 pm      User: SYSTEM
Title: Sample LU save

Verifying tape.
Verifying save of LU 12 to file 1
Verifying save of LU 57 to file 2

      Save Record
-----
File #   Tape #   LU
-----
   1       1       12
   2       1       57

ASAVE : rw
ASAVE : ll
Closed list file.
ASAVE : ex

```

ASAVE Error Messages

ASAVE error messages, error messages and warnings shared by ASAVE and ARSTR, and error messages generated by ARSTR are all presented following the discussion of ARSTR.

COPYL

COPYL

COPYL is an online utility that copies the entire contents of one disk LU to another. It is intended for users who want to make copies of disks without using FMP. COPYL directly addresses the disk driver, rather than using the FMP routines, thus increasing the speed of the copy operation. It ignores the file structure in the copy process, addressing tracks and sectors on the disk, rather than records in the individual files. The speed of the copy, therefore, is independent of the contents of the disk.

Note that the utility can only copy like media; in other words, you can copy an HP 9895 disk to another 9895 disk, but not to an HP 7908 disk. The disk need not be in FMP format; COPYL copies LIF disks, HP 125 compatible disks, or any other disk that has a physically compatible recording format.

The available memory following the last location used by the COPYL program itself becomes the buffer for the copy operation; a minimum of 128 words (equivalent to one sector) must be available. The larger the buffer area, the more efficiently COPYL operates. If the memory buffer is not large enough to copy a full sector at a time, you can increase the size of the partition for COPYL by using the loader SZ command.

See the sections “Backup Utilities” and “File Interchange on RTE-A” in Chapter 1 to learn when to use COPYL and other physical backup utilities.

Calling COPYL

Before you invoke COPYL, dismount the destination disk LU with the “DC,lu” command. You can then call COPYL interactively, or in the runstring, as follows:

```
CI> [RU, ] COPYL[ ,sourceLU ,destinationLU ]
```

The *sourceLU* and *destinationLU* are the logical unit numbers of the disk to copy and the disk to receive the data. If you do not specify the source and destination, COPYL prompts you:

```
Enter "From" LU, "To" LU <EX to Exit>
```

If you specify the source and destination improperly (that is, not as an integer), the following message appears:

```
Bad parameter format.  
Format is from, to; i.e., 32,31
```

The Copy Operation

When you enter the source and destination parameters correctly, COPYL displays the following message:

```
Contents of LU xx will be destroyed.  Enter "GO" to proceed.
```

When you enter GO, COPYL begins the copy operation, overlaying the destination disk with the contents of the source disk, sector by sector. When the copy is complete, this message displays:

```
COPYL exiting; verify that "system" disk is in place  
Enter "GO" when ready.
```

The utility does not exit until you enter GO.

COPYL Error Messages

The following error messages may be issued by COPYL:

Disk is not up and available

One of the disk LUs (source or destination) is down or busy with another program.

Disk sizes do not match

The number of tracks and the number of sectors per track must be the same for both disks.

Insufficient memory in partition. Resize program and try again.

The program uses the remainder of its partition for the copy buffer. There is not enough memory behind the program for a minimum buffer (128 words).

Not a disk LU

The device type recorded in the DVT indicates that the source or destination devices are not disks or are not the same device type.

"To" disk must be dismounted.

The destination disk must be dismounted before running COPYL. Exit the program, dismount the "to" disk and try again.

DRSTR

DRSTR

DRSTR is used to restore data from removable disk media saved using DSAVE. DRSTR restores one disk LU, a group of disk LUs, or an entire disk unit.

Calling DRSTR

Invoke DRSTR with the following runstring:

```
CI> [RU, ]DRSTR[ ;command[ ; . . . ;command ] ]
```

If you do not enter a command in the runstring, DRSTR prompts you interactively, as follows:

```
DRSTR :
```

You may separate parameters within each DRSTR command with commas or blanks, and you may combine multiple commands in the runstring or in response to the prompt, separated by semicolons (;). In the runstring, you must precede each semicolon with a backslash (\) so CI will not interpret the semicolon as a CI command delimiter. You may also use the underscore (_) character to separate parameters in the runstring; DRSTR converts them later to blanks. For example:

```
CI> drstr\;me_32\;lh\;re_4:20_7
```

Commands are executed in order, from left to right. If any command causes a serious error, execution terminates with that command, and subsequent commands are ignored.

DRSTR Commands

DRSTR uses commands that let you terminate DRSTR, list all the commands and their parameters, list header information, define a list file or device, define removable disk media, specify a restore operation, and select user error handling mode. Table 2-8 summarizes the DRSTR commands, which are described in the following sections.

Table 2-8. DRSTR Commands Summary

Commands	Description
Information Command	
HE lp	Lists all DRSTR commands and parameters
Configuration Commands	
ME dia <i>lu</i>	Defines removable disk media
RE store <i>file#:lu</i>	Specifies a restore operation
UE User Error Handling ON OFF	Selects user error handling mode
Listing Commands	
LH List Header <i>file#</i>	Lists header information from removable disk
LL List File or Device <i>file_desc lu</i>	Opens a list file or list device to which information is logged in addition to the terminal
Exit Commands	
AB ort EN d EX it	These three commands terminate DRSTR

Abort, End, and Exit (AB) (EN) (EX)

Purpose: Any one of these commands terminates DRSTR.

Syntax: AB, EN, or EX

Description:

Upon termination, all disk LUs are released and any list file or device closed.

Help (HE)

Purpose: Lists all DRSTR commands and parameters.

Syntax: HE, ??, or ?

DRSTR

List Header (LH)

Purpose: Lists header information for a file from the removable disk.

Syntax: LH[,*file #*]

file # The file number on the removable disk media.

Description:

If the file number you specify does not exist on the current removable disk, the header record of the first or last file on the disk is listed, depending upon the number specified.

If you do not specify a file number, the default is to list the header information for the next save file (or the first save file if this is the first time header information is being listed from the removable disk).

If the number is smaller than that of the first save file, the header record of the first save file is listed. If the number is greater than that of the last save file on the disk, the header record of the last save file is listed.

Note

All file numbers specified are relative to the beginning of the first DRSTR removable disk, if there is more than one. If the save file for a disk LU crosses removable disk, it has the same file number on all removable disks.

List Device (LL)

Purpose: Opens a list file or list device to which terminal displays and error messages are to be sent in addition to the terminal.

Syntax: LL[, <*file descriptor*> | *lu*]

file descriptor The list file.

lu The list device.

Description:

A defined list file is created if it does not already exist, and a list device is opened. Any previously opened list file or device is closed first. If you do not specify a file or list device, the currently defined list file or device is closed.

If you specify an existing file, a duplicate file error occurs. In this case, if user error handling is ON, an error message is issued and the LL command terminates; otherwise, the command is skipped.

If you specify a closed list file or a device that is down and user error handling is ON, DRSTR prompts you for corrective action or for a break before continuing. If user error handling is OFF, the executing command default is to break. If user error handling is OFF and a break is caused by one of the above mentioned errors, the list file or device is closed and the executing command continues to execute.

Media (ME)

Purpose: Defines the removable disk LU number used by DRSTR commands to access the removable disk.

Syntax: ME [, *lu*]

lu The logical unit number of the removable disk medium. If *lu* is omitted, either the previously defined LU or zero is displayed.

Description:

You may also use ME to examine the currently defined LU. You must enter ME before a restore operation.

The LU you specify must not be mounted. It is locked and remains locked until DRSTR terminates or you define another LU.

An error occurs if DRSTR accesses a removable disk not defined by the ME command. This typically happens when you enter an RE or LH command without or before the ME command.

To find out the LU numbers associated with the removable disks, consult the system generation answer file or check with a superuser (or your system manager).

Restore (RE)

The RE command restores data from the removable disk to the LU specified, either file-by-file for one or more files (“LU restore”), or a group of files from the same drive unit (“unit restore”). It has two formats, one for each type of restore, described separately below.

File-by-file (LU) Restore:

Purpose: Specifies a file-by-file (LU) restore.

Syntax: RE , <*file # :lu*> [, . . . , <*file # :lu*> [, VE]

file # :lu The file number on the removable disk and the disk LU where the file is to be restored.

If you specify one file number, then all the parameters must specify a file number. If you do not specify a file number, the default file is 1 for the first disk LU specified, 2 for the second, and so on.

The colon is required when entering the *:lu* parameter. If you do not enter an LU, the destination disk is determined by examining the header record of the save file for the disk LU saved. If you do specify *lu*, the parameter is specified as *file #* (no colon).

DRSTR

Description:

Only one of the above forms is allowed; combinations of file number and LU defaults result in a DRSTR error. For example, the forms

```
re, 3:20, :22
```

or

```
re, 6, :30
```

cause this message to display:

```
Error - Illegal file #:lu parameter combination.
```

The following combination is acceptable:

```
re,4:30,7
```

Note

All file numbers specified are relative to the beginning of the first removable disk, if there is more than one. If the DSAVE file for a disk LU continues on another disk, it has the same file number on both removable disks.

Unit Restore:

Purpose: To restore a group of files from the same drive unit.

Syntax: RE[, <file #>UN[, VE]]

file # The number of the first save file of the group of save files to be restored from the removable medium. If omitted, file 1 is assumed.

Description:

The restore operation only restores the specified DSAVE file and any subsequent files from the same disk unit saved in the same operation.

RE command options are described in the next section. The restore operation is explained in detail in the section “The Restore Operation” later in this chapter.

RE Command Options

The RE command options parameter lets you specify whether you want to have verification of the data when it is restored, and whether you want to specify a unit restore.

Options may be specified in any order. Table 2-9 summarizes the RE command options.

Table 2-9. RE Command Options Summary

Options	Description
VErify	Verifies transferred data when a restore operation is performed
UNit Restore	Specifies a unit restore, rather than a file-by-file restore

Verify Data (VE)

You may have transferred data verified when the restore operation is performed. To select verification, simply enter VE in the RE command option parameter, as shown above in the RE format.

Unit Restore (UN)

The UN option specifies a unit restore beginning at the file you enter, or file 1 if you do not enter a file number, as described above.

User Error Handling (UE)

Purpose: Sets the user error handling mode or displays the current mode if no parameter is entered.

Syntax: UE [, ON | OFF]

ON You are prompted for corrective action if a recoverable error occurs. The initial setting is ON.

OFF The default course of action is taken, as described in the discussion of each command for default information.

DRSTR

The Restore Operation

All disk LUs being restored are first checked for valid disk LU numbers. Checking may occur when you enter the RE command, or when each save file is encountered during the restore.

The sections that follow describe the various elements of a restore operation.

Checking the Disks

Disk LUs entered in the RE command are checked when you enter the command. An illegal disk LU causes the restore to terminate at once, and the following error message is displayed:

```
Error - LU xx not a supported disk.
```

Disk LUs read from the save files are checked when each save file is encountered during the restore. If any LU is illegal in the current system, the following warning message is displayed:

```
Warning - LU xx not a disk in current system.
```

If user error handling is ON, DRSTR prompts you to break, skip this restore, or enter the disk LU to be restored. If user error handling is OFF, the restore of this file is skipped.

Sectors per Track Must Match on Disks

DRSTR only restores a DSAVE file on removable disk to a similar disk, one with the same sectors per track configuration. If the sectors per track are different, the following error is displayed and the restore of that LU is skipped:

```
Error - Sectors per track mismatch.
```

Total Number of Tracks May Differ

The total number of tracks of the saved disk (source) need not be the same as the number of tracks on the disk being restored (destination). If there is a difference, one of the following messages is displayed:

```
Warning - Saved disk has more tracks than target LU xx.
```

or

```
Warning - Saved disk has less tracks than target LU xx.
```

After the warning message is displayed, you will see the following messages:

```
Saved disk tracks = xx  
Target disk tracks = xx
```

If user error handling is ON, you are asked:

```
Restore this disk (Y/N) ?
```

If you respond by entering “N”, the restoration of the disk is skipped. If you respond “Y”, the disk is restored. When you restore the disk, be aware that the destination LU size is not the same as the DSAVE file on the removable media. A smaller destination LU results in the restoration of only part of the DSAVE file.

A destination LU of a different size restores the save file, but if there is a file system on the LU, it is not accessible, because the file system directory information is at the last track of the LU. The directory information of the restored file system is thus not at the proper location.

If user error handling is OFF, the default is to restore the disk.

Disk Locking

Each disk LU being restored is locked when the restore operation begins, to ensure the integrity of the restored data. The disk remains locked either until the restoration is complete (if you did not select VE) or until after verification (if VE was specified).

If the destination disk cannot be locked (for example, it is currently locked to another program), and user error handling is ON, the following message is displayed:

```
Error - LU xx locked to (program name) Correct problem and type GO  
to continue, SK to skip or BR to break.
```

If user error handling is OFF, the restore of that disk is skipped.

When a disk is locked, it is checked to see if it is mounted; if it is, the following message is displayed:

```
Error - LU xx is mounted.
```

If user error handling is ON, DRSTR prompts you as follows:

```
Correct problem and type GO to continue, SK to skip or BR to break.
```

If user error handling is OFF, the restore of that disk is skipped.

Order of Disk Restoration

Disks are restored in the order in which they were saved to removable disk media, not as they are specified in the RE command.

DRSTR

Unit Restore

For a unit restore, you must first enter the removable disk that contains the first section of the save file specified in the RE command; otherwise, DRSTR prompts you with one the following messages:

```
Error - First file is x section y - searching for file z.
```

or

```
Error - Last file is x section y - searching for file z.
```

DRSTR then prompts you as follows:

```
Correct problem and type GO to continue or BR to break.
```

Note

If the next save file to be restored from is NOT on the current removable disk, you can insert the removable disk volume that contains the first section of the save file. (You may skip over the removable disk volumes that will not be used during this restore operation.)

Restoring Fixed Disks from Removable Disks

Each removable disk must be a valid, DSAVE format removable disk. If it is not, the following message is displayed:

```
Error - Non-DSAVE removable disk media.
```

```
Insert another removable disk media and type GO to continue or BR to break.
```

If the save file to be restored is not on the removable disk, one of the following messages is displayed:

```
Error - First file is m section n - searching for file x section y.
```

or

```
Error - First file is m section n - searching for file x section y.
```

After the appropriate error message is displayed, DRSTR prompts you as follows:

```
Correct the problem and type GO to continue, SK to skip or BR to break.
```

Header Records

Information for the restore of a disk is obtained from the system and the header record. Data from the header record is used to control the restore. If a checksum error occurs on the header record, the restore of that disk is skipped and you are asked for another removable disk.

Because the location of the next header record on the removable disk is in the header record, any save file(s) after the current one on this removable disk media must be skipped. When this occurs, the following error message is displayed:

```
Error - Header record checksum.  
Insert another removable disk media and type GO to continue or BR to break.
```

When the actual restore of the disk LU begins, information from the header record of the save file is logged.

Restoring the Data

For each save file, disk data from the removable disk is restored to the disk LU until the destination disk is full or the end of the disk data in this section of the save file is reached. If an error occurs while reading a part of the data record from the removable disk during the restore, the following error message is displayed:

```
Error - Data record read error.
```

The range of disk block numbers is determined and reported in the form of the following message:

```
Data targeted for disk blocks x to y.  
Continuing.
```

In any case, the restore continues. When the end of this section of the save file is reached, or the target disk LU is full, the disk data is verified (if you selected VE). If the restoration is not complete (that is, the target LU is not full and this was not the last section of the save file), DRSTR prompts:

```
Insert next removable disk and type GO to continue or BR to break.
```

If the restoration of the target LU is complete, the next save file to be restored from is searched for on the current removable disk. When it is found, the restore begins. If it is not found, DRSTR prompts you as follows:

```
Insert next removable disk and type GO to continue or BR to break.
```

If the removable disk drive or destination disk goes down during the restore, the following message is displayed:

```
LU xx is down.  
Correct problem and UP device, or use BR to break.
```

You may also see system messages about the device that is down.

DRSTR

When the restore of a disk LU completes, the following message notifies you:

```
Restore of file n to disk LU xx complete.
```

Verification

When you select VE, the restore operation is a two pass process. During the first pass (restore), data is read from the removable disk and written to the fixed disk. In the second pass (verify), the data is re-read from removable disk into one buffer, and the corresponding data restored to fixed disk is read into a second buffer. The data in the two buffers are compared bit-by-bit.

If a mismatch between the two buffers occurs, the following message displays, and verification continues to the end of the disk or until three verify errors occur on that disk restore, causing the process to abort:

```
Warning - Possible data loss in disk block range n to n.
```

Break Detection

If DRSTR detects a break, it terminates all ongoing operations and unlocks all LUs except the removable disk drive LU.

Loading DRSTR

Load DRSTR online, using the LINK loader and the LINK command file #DRSTR. The DRSTR software and the command file are included in the 92077A RTE-A Master Software. DRSTR requires 17 pages of memory (including an extra 250 words for buffering). The recommended program size is 32 pages.

DSAVE and DRSTR Error Messages

The sections that follow list 1) all the messages that DSAVE generates, 2) error messages that are shared by DSAVE and DRSTR, and 3) the error messages that DRSTR generates.

DSAVE Error Messages

Error – Cannot overwrite removable disk media in the current media set.

The removable disk media has a valid DSAVE format and is from the same set of the current save operation.

Error – Data record verify failed.

Retry the operation. If the error persists, try another media or disk drive unit.

Error – Illegal unit save.

You entered two or more LUs in a unit save operation. Only one LU is required.

Error – Media LU must not be in the list of LUs to be saved.

The media LU was specified as one of the LUs to be saved. Retry the operation and specify the correct LUs.

Error – No disks to save.

The LU parameter was missing from the SA command.

Error – Only *xx* save file number available.

The save file number specified is greater than the number of save files available on the removable media.

Error – Too many LUs specified.

Only 63 LUs are allowed in the save operation. Save the excess LUs in another operation.

Error – Unknown command; use ?? for help.

This is probably due to a typing error. Reenter the correct command.

DRSTR

DSAVE and DRSTR Shared Error Messages

Error – Bad header record detected.

Header record verification failed. Retry operation.

Error – Bad LU.

Illegal LU specified. For example, the LU specified in the LL command must be an I/O device LU.

Error – Disk retries exhausted.

Unsuccessful disk access. Retry the operation with another disk unit.

Error – File number expected.

A file number must be specified in the LH command.

Error – <FMP error message>.

A file system error occurred while accessing the list file or device.

Error – Header record verify failed.

DSAVE or DRSTR could not verify the header record of an LU, perhaps because of bad media. Use the LH command to check the header in question. If an invalid header is found, this LU and the one after it are not recoverable. To save them, repeat the save operation for those LUs.

Error – Illegal removable disk media LU.

The LU specified in the ME command is not a removable disk LU or is not found in the system. Check the system LU assignment and reenter the ME command.

Error – IOxx on LU xx <error description>

A device I/O error was detected. The I/O error messages can be found in the *RTE-A Quick Reference Guide*, part number 92077-90020, or the *RTE-A User's Manual*.

Error – LU xx is a mirrored volume LU.

The physical paired disk LU associated with the mirrored volume LU must be specified, instead of the mirrored volume LU. (Refer to the *Data Pair/1000 Reference Manual* for more information.)

Error – LU xx is down.

The specified I/O device is down. Ensure that the device is operational, and bring it up.

Error – LU xx is locked to <program name>.

The specified LU is locked by a program. That program must be terminated before reentering a command.

Error – LU *xx* is mounted.

Dismount the LU to be saved or restored, then retry the operation.

Error – LU *xx* not a supported disk.

The specified LU is not a supported disk in the RTE-A system. Check the system configuration for the proper LU number.

Error – Non-DSAVE removable disk media.

Information contained in the removable disk media is not in DSAVE format. Saving LUs on that media will destroy the existing information.

Error – Not enough memory (size program).

DSAVE or DRSTR requires at least an extra page of memory. Size up the program one page larger either temporarily, with the SZ command from CI, or permanently, with Link.

Error – Option expected.

Usually a typing error. Anything entered after an option must be an option.

Error – Removable disk media must not be a mirrored volume LU.

The LU specified in the ME command is a mirrored volume LU. Check the system LU assignment, and reenter the ME command. (Refer to the *Data Pair/1000 Reference Manual* for more information.)

Error – Removable disk media LU undefined.

A save or restore operation initiated before a media LU was defined. Specify a removable disk media LU with the ME command, and retry the operation.

Error – Removable disk media LU = *xx*.

The removable disk media LU number is specified in other commands, for example, in the LL command as the list device. Check the system configuration and redefine the LUs.

DRSTR

DRSTR Error Messages

Error – A duplicate save file number detected.

Typically, the same file number was specified, such as “re , 5 : 20 , 6 , 5”.

Error – Bad LU specification.

This error usually occurs when non-numeric characters are entered in the LU parameter field. Reenter the command with the proper LU number.

Error – Bad save file number specification.

Typically, a typing error; for example, “re , 5xx” or “re , t5”.

Error – Data record read error.

The verify was unsuccessful during the restore. Retry the operation.

Error – Illegal file #/lu parameter combination.

In a file-by-file restore, the file#/lu parameters must be consistent in all parameter specifications. They must be in one of the following forms: : <file #>, < : lu > or <file# : lu >.

Error – Illegal unit restore.

Usually, this means the LU parameter is missing from a unit restore operation.

Error – LU expected after “:”.

The LU parameter is missing from a multi-file restoration.

Error – LU must be greater than zero.

Probably non-numeric characters were entered instead of numbers. Enter the correct LU number.

Error – Media LU must not be one of the disk LUs to be restored.

Check the system configuration, and enter the correct LU numbers.

Error – No LU or save file number specified.

RE command parameters are missing. Refer to the discussion of the RE command earlier in this chapter for the correct command syntax.

Error – Save file number must be greater than zero.

This error is typically caused by a typing mistake. Reenter the command, and specify the proper file number.

Error – Save of disk LU was skipped during save operation.

The LU you specified cannot be found in the removable media.

Error – Sectors per track mismatch.

The source and destination disks have a different number of sectors per track. The restore operation is performed, but the results are uncertain. If the destination disk has a valid file system, it is no longer accessible to the system.

Error – Too many save file/lu parameters specified.

You may enter a maximum of 63 file #/lu parameters. However, the command string is subject to the file system command string length restrictions. Refer to the *RTE-A User's Manual* for details.

Error – Unknown parameter.

Probably a typing mistake was made, or the wrong command syntax entered. Refer to the DRSTR command description in this chapter for the correct command syntax.

Error – Wrong removable disk media set.

A removable disk from the wrong set was inserted. Remove it, and insert the correct disk.

DSAVE

DSAVE

DSAVE is used to save data from fixed disks to removable disks (for example, 3-1/2, 5-1/4, and 8-inch floppies, and HP 7906H removable platters). DSAVE lets you back up large, contiguous areas of disks to removable disk media regardless of the content of the disks (such as files). You may use DSAVE to save one disk LU, a group of disk LUs, or an entire disk unit to removable disk.

Because DSAVE is used to save data in one or more LUs, file structures are ignored. To save files, use the CI CO command, or use the TF utility with other media.

The size of LUs saved may be a factor when the removable media are micro or mini floppies. Large LUs may take longer to complete, and require many disks. It takes approximately 30 seconds to save without verification on a micro floppy that holds approximately 1055 blocks of data (including one block reserved for header record). With verification, the time is one minute. If you are using old, removable disks (or any used disk of uncertain integrity), verification is recommended even though the save operation takes longer.

Calling DSAVE

Invoke DSAVE with the following runstring:

```
CI> [RU, ]DSAVE[ ;command[ ;command[ . . . ] ] ]
```

Alternatively, you may run DSAVE interactively, entering commands at the following prompt:

```
DSAVE :
```

Separate command parameters either with commas or blanks. You may enter multiple commands either in the runstring or in response to the prompt, separated by semicolons. In the runstring, semicolons must be preceded by a backslash (\) so that CI does not interpret them as CI command delimiters. The underscore (_) character is also used in the runstring as a parameter separator. For example,

```
CI> dsave\ ;me32\ ;lh\ ;sa_20
```

Commands are executed in order, from left to right. If any command causes a serious error, execution terminates with that command (all the preceding commands are executed).

DSAVE Commands

DSAVE commands let you terminate DSAVE, display all the commands and their parameters, list header information, define a list file or list device, define removable disk media, specify a save operation, specify a save title, and define the error handling mode. Table 2-10 summarizes the commands, which are described in the following sections.

Table 2-10. DSAVE Commands Summary

Commands		Description
Information Command		
HElp		Lists all DSAVE commands and parameters
Configuration Commands		
ME dia	<i>lu</i>	Defines the removable disk LU
SA ve	<i>lu</i>	Specifies a save operation from fixed to removable disk
Ti tle	<i>text</i>	Sets the title to be used in subsequent save operations
UE User Error Handling	ON OFF	Sets user error handling mode and defines how user errors are handled
Listing Commands		
LH List Header	<i>file #</i>	Lists header information from removable disk
LL List Device	<i>file descriptor lu</i>	Defines a list file or device to which error messages are sent in addition to the terminal
Exit Commands		
AB ort EN d EX it		These three commands terminate DSAVE

DSAVE

Abort, End, and Exit (AB) (EN) (EX)

Purpose: Any one of these commands terminates DSAVE.

Syntax: AB, EN, or EX

Description:

When the program terminates, the list file or device specified is closed, and all locked LUs are released.

Help (HE)

Purpose: Displays a list of all DSAVE commands and the parameters of each command.

Syntax: HE, ??, or ?

List Header (LH)

Purpose: Lists the header information for a file from the removable disk.

Syntax: LH[, *file#*]

file# The file number on the removable disk.

Description:

If the specified file does not exist on the current removable disk, the header record of the first or last file on the removable disk is listed, depending upon the value of the specified file. If you do not enter a file number, the program lists the header information for the next save file (or the first save file, if this is the first time header information is being listed from the removable disk).

If the file number is smaller than the file number of the first save file on the removable disk, the header record of the first save file is listed. If the number is greater than the last save file, the header record of the last save file is listed.

Note

All file numbers specified are relative to the beginning of the first DSAVE removable disk, if there is more than one. If the save file for a disk LU occupies more than one removable disk, it has the same file number on all the disks.

Table 2-11 illustrates LH command use, followed by descriptions of the various fields in the header.

Table 2-11. LH Command Example

DSAVE : LH,1

Created Using: DSAVE 92077-16702 REV.6000 <920903.1642>
 Volume Set Date: Wed Dec 9, 1992 3:06:41 pm
 Volume #: 1 File #: 1 Options: Disk LU: 61
 Tracks: 6 SecTrk: 62 Section: 1 Trk: 0 Sec: 0
 Created: Wed Dec 9, 1992 3:06:46 pm User: SYSTEM
 Title: Sample header

Created Using	<i>Information on the version of DSAVE used to save the disk.</i>
Volume Set Date	<i>Time stamp of creation of the first file, used to identify a multi-volume set.</i>
Volume#	<i>The number that identifies a set of removable disks.</i>
File#	<i>Save file number.</i>
Options	<i>Options selected for the save operation.</i>
Disk LU	<i>Disk LU saved.</i>
Tracks	<i>Total number of tracks for the disk LU.</i>
SecTrk	<i>Sectors per track for the disk LU (64-word sectors).</i>
Section	<i>Save file section number, used when a save operation crosses disks.</i>
Trk	<i>Disk track where disk data in this section of the save file starts.</i>
Sec	<i>The sector in track described above.</i>
Created	<i>Time stamp when disk was saved.</i>
User	<i>Name of the user account from which DSAVE was scheduled.</i>
Title	<i>Title specified for the save operation.</i>

DSAVE

List Device (LL)

Purpose: Defines a list file or list device to which error messages and terminal displays are to be sent in addition to the terminal.

Syntax: LL [, <file descriptor> | *lu*]

file descriptor The file to be opened.

lu The list device.

Description:

If the list file you define does not already exist, it is created, and a list device is opened. (A previously opened list file or device is closed before the new list file is created or the list device is opened.) If you do not specify a file descriptor/LU, the currently defined list file or device is closed.

A duplicate file error occurs if the file you specify already exists. If user error handling is turned ON, the LL command terminates with an error; otherwise, the command is skipped.

If you specify a device that is down or a closed list file, you are prompted to correct the problem or enter breakmode if user error handling is ON. If user error handling is OFF, the default is to break. In this case, the list file or device is closed, and the executing command continues to execute.

Media (ME)

Purpose: Defines the removable disk LU.

Syntax: ME [, *lu*]

lu The logical unit number of the removable disk. If omitted, either the previously defined LU or zero is displayed.

Description:

This LU number is used by DSAVE to access the removable disk. You can also use ME to examine the LU currently defined.

The specified LU must not be mounted. It is locked and remains locked until DSAVE terminates or another LU is defined. You must enter ME before the SA and LH commands; otherwise, the message, "Removable disk media LU undefined" is displayed.

To find out the range of LUs associated with the removable disks, consult the system generation answer file, or check with a superuser (or your system manager).

Note The removable disk media LU must be generated into the operating system, and the entire removable media must be a single LU; otherwise, there may be problems. For example, the two surfaces of an HP 7906 removable platter must be generated in the system as one LU in cylinder mode.

Save (SA)

Purpose: Specifies a save operation from a fixed disk to a removable disk.

Syntax: `SA, lu [, lu [, . . . , lu [, VE [, NL]]]]`

or

`SA, lu, UN [, VE [, NL]]`

lu The logical unit of the disk.

Description:

You may perform a save operation on one contiguous area specified by one LU, or as a unit save of the whole drive unit.

The save operation is described in detail in the section “The Save Operation” which follows the discussion of the remaining DSAVE commands. The SA command options are described below.

SA Command Options

The SA command options parameter lets you specify one or a combination of the options shown below. Options may be specified in any order. Table 2-12 summarizes the options.

Table 2-12. SA Commands Option Summary

Options	Description
No Locking	Disk LUs are not locked during a save operation
UN it Save	Specifies a unit save of the disk pointed to by the specified LU
VE rify	Instructs DSAVE to verify the save operation

DSAVE

No Locking (NL)

NL specifies that the disk LUs are not to be locked during the save operation. The default is to lock them during the save to ensure the integrity of the saved disk data.

If you select both the NL and VE options with the SA command, errors may occur during verification if disk data was altered (written) between the save and verify passes while the disk was unlocked.

Unit Save (UN)

UN specifies a unit save of the disk unit pointed to by the specified LU. If you do not specify UN, the save is considered to be an LU save of all the LUs specified.

Verify (VE)

The VE option instructs DSAVE to verify the save operation. The verification method is described later in the section called “Verification”. Verification is recommended to ensure a successful save operation, as it avoids problems such as bad disk media or disk drive errors. If you use old removable disks for a save, VE is highly recommended.

Title (TI)

Purpose: Sets the title to use in subsequent save operations.

Syntax: TI[, *text*]

text The text for the title, up to 40 characters.

Description:

The specified title is placed in the header record of the save files on the removable disk. Any characters over 40 are truncated. If you do not enter any text, the title consists of blanks.

When you enter TI in the runstring, be aware that CI interprets both blanks and commas as parameter delimiters. If you do not want to enter any text for the title, you may enter the underscore (`_`) character to replace the default blanks; DSAVE later converts the underscore(s) to blanks.

CI also converts all lowercase characters to uppercase in the runstring before DSAVE retrieves it. For example:

```
CI> dsave\;ti Unit_save_of_LU_37_without_VE
DSAVE : ti
Title = UNIT SAVE OF LU 37 WITHOUT VE
DSAVE :
```

However, when you enter the title interactively, it remains the same as typed. For example:

```
DSAVE : ti Unit save of LU 37 without VE
Title = Unit save of LU 37 without VE
DSAVE :
```

User Error Handling (UE)

Purpose: Sets the user error handling mode and defines how user errors are handled.

Syntax: UE[,ON|OFF]

ON	You are prompted for corrective action if a recoverable error occurs. The initial setting is ON.
OFF	The default action is taken, as described in the following sections.

Description:

If you correct the error, command execution continues.

You may also use UE to see whether user error handling mode is ON or OFF by entering the command without any parameter. For example:

```
DSAVE : ue                                (Displays current UE mode)
User Error handling = ON
DSAVE : ue off                             (Turns UE off)
User Error handling = OFF
DSAVE :
```

The Save Operation

Before the save operation begins, all the disk LUs to be saved are verified as legal disk LUs. If any LU is illegal, DSAVE displays the following error message:

```
Error - LU xx not a supported disk.
```

If you are running DSAVE interactively, you will be prompted with additional messages; otherwise, it terminates.

If either the removable disk or the source disk LU goes down during the save operation, DSAVE issues the following message:

```
LU xx is down.
Correct problem and UP device or use BR to break.
```

System error messages about the downed device may also display. The sections that follow describe the various elements of a save operation.

DSAVE

Unit Save LU Checking

If the SA command specifies a unit save operation, DSAVE locates all the disk LUs associated with the disk unit pointed to by the disk LU you entered, and proceeds in the same way as an LU save.

Disk Locking

If you do not specify NL for the save, each disk LU being saved is locked just before it is saved, and remains locked until the save completes unless you specified VE, when it remains locked until after the save is verified.

If you do not specify NL, and the disk to be saved cannot be locked (it is currently locked to another program), if user error handling is ON, you are prompted with the following message:

```
Error - LU xx locked to program name
Correct problem and type GO to continue, SK to skip or BR to break.
```

If user error handling is OFF, the SA command is skipped and the disk LU is not saved.

Removable Disk Overwriting

At the beginning of a save operation, the removable disk in the media specified is checked for a valid file system disk header or a DSAVE format media. When a valid CI disk volume header or FMGR disk cartridge header is found, the following message is displayed:

```
Warning - Removable disk media has a valid file system disk header.
Do you want to overwrite the media (Y/N) ?
```

If you respond “Y”, the removable disk to be used in the save operation is initialized. Information on that disk stored in CI or FMGR file format is destroyed. If you respond “N”, DSAVE prompts for another removable disk media with the following message:

```
Insert another removable disk media and type GO to continue or BR to break.
```

The removable disk is checked for a valid DSAVE format (that is, it does not have a valid file system disk header, and it has a valid DSAVE header record for the first save file on the removable disk). If a valid DSAVE format is present and the disk is from the same media set of the current save operation, DSAVE issues the following message:

```
Error - Cannot overwrite removable disk media in the current media set.
Insert another removable disk media and type GO to continue or BR to break.
```

If the disk has a valid DSAVE format but is from a different media set, the information from the first header record on the removable disk is displayed, along with this message:

```
Warning - Removable disk media is a DSAVE format media.
Do you want to overwrite the media (Y/N) ?
```

When you enter “Y”, the media is used in the save operation. If you enter “N”, the following message appears:

```
Insert another removable disk media and type GO to continue or BR to break.
```

If you enter BReak, the save operation terminates and you are returned to the DSAVE prompt.

Save File Records

DSAVE files contain two types of records, header and data. A header record at the beginning of each save file contains information about the save. The user information included in the header record is shown below:

- Version of DSAVE used
- Time stamp of save operation
- File number (relative to beginning of the media set)
- Save options specified
- LU of disk area saved
- Number of tracks of the LU saved
- Section number (used when a file is saved on more than one removable disk)
- Status (end of save, continues on next removable disk with new fixed disk LU, or continues on next removable disk with current fixed disk LU)

The data record is a variable length record that contains physical images of the portion of the disk being saved. A more detailed explanation of these records and the DSAVE removable disk media format is provided in the section on “DSAVE Removable disk Media Format” later in this chapter.

Saving to the Removable Disk

When a save operation begins, the disk is locked unless you selected the NL option. disk data is read from the disk LU and written to the removable disk until the end of a disk LU is reached or until the removable disk is full. The header record is then written.

If the removable disk is full, the portion of the save operation that is on the removable disk is verified if you selected the VE option. Then the following message is issued:

```
Insert next removable disk media and type GO to continue or BR to break.
```

The header for the last section of the last save file is marked with the “end-of-save operation and end-of-media” flag before it is written to the removable disk media. This completes the save operation.

DSAVE

Verification

When you select the VE option to use with the SA command, the save operation becomes a two pass process. During the first pass (save), disk LUs are saved to the removable disk. In the second pass (verify), the removable disk is repositioned to where the save began on the removable disk, and the following messages are displayed (the second message is displayed for each save file on the media.):

```
Verifying removable disk media.  
Verifying save of LU xx to file x.
```

For each save file, the checksum of the header record is verified. If a checksum error occurs while verifying the header record, the following message is displayed:

```
Error - Header record verify failed.
```

The save operation continues, but you should use the LH command to check the header records of the LU in question and the one immediately after it. If the header records are corrupt, repeat the save operation for these LUs.

The data records are verified by reading the data records from removable disk into one buffer, and the corresponding data on the source disk LU into a second buffer. The two buffers are then compared, bit-by-bit. If a data record comparison fails, the following is displayed:

```
Error - Data record verify failed.
```

If a verify error occurs on a save file, the appropriate error displays and the rest of the section of the file is skipped. The verify continues with the next file.

Break Detection

DSAVE detects a break whenever you enter the BR command in response to a prompt which has BR as one of the responses, or when you enter a CI, CM, or FMGR break command. When a break occurs, DSAVE terminates the current save operation, unlocks all LUs that were locked (except the removable disk LU), and attempts to clean up the save operation.

If DSAVE is filling a removable disk media with save files when it detects the break, DSAVE stops copying disk data from the disk LU being saved and writes the header record with a status that indicates the save operation was broken and the save file is incomplete.

If the media LU is down, the following message is displayed:

```
Warning - Volume is corrupt.
```

If DSAVE cannot clean up (for example, the removable disk media LU is down, or you aborted DSAVE), a corrupt DSAVE media and media set result. The removable disk is corrupt because the last header on the media is NOT marked as end of media set.

If DSAVE is not filling a removable disk media when it detects a break (for example, when the volume is being verified or when prompting for another media), DSAVE takes the following steps:

1. Checks if the media LU is down; if so, DSAVE displays:

```
Warning - Media set corrupt.
```

2. Checks whether the media is a DSAVE format media; if not, DSAVE displays:

```
Error - Non#DSAVE removable disk media.
```

3. Checks if the media is from the current media set; if not, DSAVE displays:

```
Error - Removable disk media is from the wrong media set.
```

4. Checks if the media is the last volume written; if not, DSAVE displays:

```
Error - Removable disk media is volume xx.
```

In case of an error, the following message is displayed after the error message:

```
Insert volume xx of the current media set and type GO to continue or BR to break.
```

When you enter GO, the program repeats steps 1 through 4 for the new media. If you enter a break, the following message displays:

```
Warning - Media set corrupt.
```

In this case, if DSAVE cannot clean up, the result is a corrupt DSAVE media set, because the status of the last header record on the last DSAVE removable disk media of the set was not marked with “end of media set”.

If the last save file on the volume was to be continued on the next removable disk media, it is corrupt because the status of the header record was not marked “save of disk LU incomplete.”

Loading DSAVE


Use the LINK loader and the LINK command file #DSAVE to load DSAVE online. The DSAVE software and the LINK command file are included in the RTE-A Master Software. No special requirements are needed. DSAVE needs at least 17 pages of memory (including an extra 250 words for buffering). The recommended size is 32 pages.

DSAVE


DSAVE Removable Disk Media and File Formats

Removable Disk Media Format

The DSAVE removable disk media format for a single LU save is shown below.




A unit or multiple LU save is a series of single LU saves. A unit save starts with the lowest numbered LU on the unit and proceeds through the highest numbered LU on the unit. In a multiple LU save, the LUs are saved in the order in which they were specified in the SA command. The DSAVE removable disk media format for unit and multiple LU saves is shown below.



Removable Disk Media Records

The DSAVE file has a fixed length header record followed by a variable length data record.



Header Records

The format of the header records is shown below.

Word	Content
0-1	Checksum
2	Length of header record
3	Program revision number (numeric)
4-10	Time stamp (ASCII, <000000.0000>)
11-13	Backup utility program name (DSAVE)
14-16	Media set ID (time stamp)
17	Volume number of media
18	Save file number (relative to beginning of media set)
19	Disk LU saved
20	Numeric device type of the disk LU
21	Select code of the disk LU
22	Device address of the disk LU

Word	Content
23	Unit number of the disk LU
24	Volume number of the disk LU
25	Cylinder starting
26	Head block
27	#Surfaces for CS/80
28	#Tracks
29	#Spares
30	#64 word sectors/track
31	Reserved
32	Save file section number
33–34	Disk block number where data for this section starts
35–37	Time stamp when this LU was saved
38	Options selected for the save of the disk LU(s)
39–46	User account name where backup was scheduled
47–66	Title specified for the save operation
67	Total number of tracks for the media LU
68	Sectors per track for the media LU
69–71	Time stamp when the save operation began
72	Header status word (described below)
73–74	Length of the disk data in this section
75	Saved disk data word from “directory”
76–77	Block address of the next header record
78–79	Block address of the previous header record
80	First save file number in save operation
81	Last save file number in save operation
82	First volume number of save file
83	Last volume number of save file (calculated)
84–127	Unused (reserved)

Header record status word:

15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
EOS	LSK	LIN	SBR	r	r	r	r	r	r	r	r				section status

- EOS – End of the media set.
- LSK – The save of the disk LU was skipped.
- LIN – The save of the disk LU is incomplete.
- SBR – The save operation was broken.
- r – reserved (zero).

DSAVE

Section status:

- 1 – Save operation continues on the next volume with a new disk LU.
- 2 – Save operation continues on the next volume with the current disk LU.
- 3 – Save operation continues on the current volume with a new disk LU.

DSAVE Error Messages

DSAVE error messages, error messages shared by DSAVE and DRSTR, and error messages generated by DRSTR are all presented following the discussion of DRSTR.

ERTSH

ERTSH is a disk exerciser utility program that tests the functionality of a disk and its controller, including the disk media and the mechanisms needed to read and write data. It can be used with any of the following disks connected via the HP 12022A disk controller:

- 10MB Winchester, Seagate ST-412
- 15MB Winchester, Seagate ST-419
- 20MB Winchester, Seagate ST-225

The ERTSH utility is generally used with FORMA to find and verify faulty disk sectors so they can be spared. Note that FORMA and ERTSH require at least 256 KBytes to run.

You can perform a series of reads from the disk without erasing data on it, in order to test the entire disk unit or a portion of the disk. A write/read test is more thorough, because more of the disk and disk controller are exercised; however, a write/read test erases data by writing to, then reading from the disk. A variety of test patterns are used in ERTSH's write/read test, thereby increasing the likelihood of detecting errors with the disk media or mechanism.

The number of test passes, error list options, and other user definable options make ERTSH a useful tool in many situations.

ERTSH Structure

A standard test pass of ERTSH consists of two or three phases. First is the standard read or write/read phase, which uses large disk requests to detect as many problems as possible in a short amount of time. This phase terminates when the desired number of passes have completed or the BR (Break) command is used.

For write-then-read tests, a cleanup phase is included so that the disk is left in a consistent state. ERTSH does not respond to a BR command when cleaning up.

To run ERTSH, enter the following runstring:

```
CI> RU,ERTSH[ ,commands ][options ]
```

Separate commands that you enter in the runstring with colons. Use blanks or commas to separate command parameters.

Alternatively, you can run ERTSH interactively and enter commands and parameters when you are prompted. If you do not enter commands in the runstring or an EX (Exit) command, ERTSH enters interactive mode.

For example, the runstring

```
CI> ru,ertsh lu 15:wr
```

executes a write-then-read test on disk LU 15, then enters interactive mode. If you include an EX command in the runstring,

```
CI> ru,ertsh lu 15:wr:ex
```

ERTSH

ERTSH terminates after the write-then-read test.

If ERTSH detects an error in a command while processing a string of commands, it displays an error message and ignores the rest of the commands. In interactive mode, ERTSH returns with another prompt or terminates if it is processing runstring commands.

ERTSH Commands

The commands available for use with ERTSH are summarized in Table 2-13 and discussed in the following sections.

Table 2-13. ERTSH Commands Summary

Commands		Description
Information Command		
HElp	??	Displays all ERTSH commands
Testing Commands		
BP Build Test Pattern	<i>patterns</i>	Defines or displays the user definable test data pattern.
CY Define Cylinder Range	<i>low/high</i>	Defines a subrange of disk cylinders for error rate test.
HD Define Disk Head Range	<i>low/high</i>	Defines subrange of disk heads for ERT test.
LL Extra Log	<i>lu</i>	Displays the current extra log device or sets up an extra log device.
LU Select disk LU to Test	<i>lu</i>	Selects the disk LU for testing.
P# Set Test Pass	<i>pass count</i>	Defines the number of passes over the disk range for an error rate test.
RO Execute Read-Only	SQ RN	Executes a read-only error rate test over the specified disk range.
SC Define Disk Sector	<i>low/high</i>	Defines a subrange of disk sectors for error rate test.
SS Spare Sectors	HD AL	Spares all defective sectors found during a test pass.
TP Select Test Pattern	<i>pattern num</i>	Selects one or more of the 19 defined test data patterns for the write-then-read error rate test.
WR Execute Write/Read Error Rate Test	SQ RN	Executes a write-then-read error rate test over the specified disk range.
Exit Command		
EX Exit ERTSH		Exits the ERTSH utility

Help (??)

Purpose: Displays a list of all the available commands or specific information about the command entered.

Syntax: ?? [*command*]

Description:

If you enter an invalid command, a list of all the commands is displayed.

Build Test Pattern (BP)

Purpose: Defines or displays the user-definable test data pattern.

Syntax: BP [*word*] ... [*word*]

word The test data pattern, from one to six words long.

Description:

Each word of the pattern can be decimal, octal, or ASCII data. For example, you can specify a test pattern of all ASCII blanks as 20040B. Since blanks, commas, and semicolons have special meaning in the command syntax, you may not use them as ASCII values.

When you enter the BP command without any parameters, the program displays the test patterns and shows which ones are currently in use.

The following example shows the patterns listed with BP:

```
ERTSH (?? for help) > bp

Active test patterns = 1 2 3 4 5 19

19) 177777B
```

Define Cylinder Range (CY)

Purpose: Defines a subrange of disk cylinders for error rate test.

Syntax: CY [*low cylinder*] [*high cylinder*]

low/high cylinder The cylinder range.

Description:

You can define a sub-area of a disk LU by specifying a cylinder range, a head range, and a sector range. The default cylinder range is all cylinders within the current test LU. You must define a test LU with an LU command prior to setting the cylinder range.

ERTSH

If you enter only one value, the high and low cylinder are set to that value. If you do not enter any parameters, the current cylinder range and the maximum cylinder range are given as follows:

```
Current cylinder range is 120 to 125
Maximum cylinder range is 104 to 158
```

EXIT (EX)

Purpose: Exits ERTSH after any necessary cleanup is performed.

Syntax: EX or /E

Description:

ERTSH does not respond to a BR (Break) command during the cleanup process.

Define Disk Head Range (HD)

Purpose: Defines a subrange of disk heads for ERT test.

Syntax: HD [*low head*] [*high head*]

low/high head The disk head range.

Description:

You can define a sub-area of a disk LU by specifying a cylinder range, a head range, and a sector range. The default head range is all disk heads within the current test LU. You must define a test LU with the LU command prior to setting the disk head range.

If you enter only one parameter, the head range is defined as that head. If you enter no parameters, the current head range and the maximum head range are given as follows:

```
Current head range is 2 to 3
Maximum head range is 0 to 3
```

Extra Log (LL)

Purpose: Displays the current extra log device or sets up an extra log device.

Syntax: LL [*lu*]

lu The LU of the log device.

Description:

All user messages are written to the extra list device if it exists. Setting the LU to zero terminates the use of the extra log device. Entering the LL command with no parameters causes the current extra log device to be displayed.

If the extra log device is not ready, or an error occurs while writing to the log device, ERTSH discontinues using the device and issues an error message.

Selecting Disk LU to Test (LU)

Purpose: Selects the disk LU for testing.

Syntax: LU [*lu*]

lu The disk test LU number.

Description:

The LU does not have to be mounted to the file system. It cannot be mounted if a write-then-read error rate test is to be performed, because all disk data is overwritten. You must define a disk test LU before a CY, HD or SC command can change the test range or before an RO or WR command can execute a test.

If you do not specify an LU, a list of all disk LUs in the system that can be tested is displayed, along with the currently selected LU.

The following is a sample list of disk LUs:

```
Sel. Code  Unit  LUs
      32B      0   10, 15, 36, 37, 38, 39, 55, 56, 57, 58, 59, 60
```

```
Current disk test LU is 15
```

Set Test Pass Count (P#)

Purpose: Defines the number of passes over the disk range for an error rate test.

Syntax: P# [*number*]

number The number of passes, from 0 to 32767.

Description:

Entering a zero as the number parameter specifies an infinite number of passes.

For test passes that use sequential access, one pass is defined as the number of disk accesses required to check each part of the disk test range once. Because random access tests are not guaranteed to access the whole disk range, one (random) pass is defined as the number of disk accesses required to randomly reach a defective track in the middle of the test disk range.

The pass count has an initial value of 12 so that a test can be performed without prior execution of this command.

If you enter the P# command without a parameter, the current pass count is returned.

ERTSH

Execute Read-Only Error Rate Test (RO)

Purpose: Executes a read-only error rate test over the specified disk range.

Syntax: RO [SQ|RN]

SQ|RN The specification of a sequential (sq) or random (rn) pass test.

Description:

The sequential access test covers the entire disk range in a uniform manner. The random access test accesses random disk locations and may not access 100% of the disk media. If you do not enter any parameter, sequential access is used.

Note that you must define a disk test LU prior to executing a read-only test.

Refer to the subsequent section on “Using ERTSH” for more information concerning the use of each type of test.

Define Disk Sector Range (SC)

Purpose: Defines a subrange of disk sectors for error rate test.

Syntax: SC [*low sector*] [*high sector*]

low/high sector The specification of the sector range.

Description:

You can define a sub-area of a disk LU by specifying a cylinder range, a head range, and a sector range. The default sector range is all disk sectors. You must define a disk test LU before a sector range is defined.

If you only specify one parameter, the sector range consists of the given sector. If you specify no parameters, the current sector range and the maximum sector range are given as follows:

```
Current Sector Range is 10 to 20
Maximum Sector Range is 0 to 30
```

Spare Sectors (SS)

Purpose: Spares all defective sectors found during a test pass.

Syntax: SS [HD|AL]

HD|AL HD only spares those sectors that contain defects that cannot be corrected. AL causes all sectors with defects to be spared.

Description:

Errors that cannot be corrected are described as error type 2, 3, or 4 in the error summary. ERTSH schedules FORMA to perform the sparing function, so FORMA must be RP'd prior to running ERTSH. ERTSH cannot spare multiple defective sectors per track or spare when the sector is unknown.

Select Test Pattern (TP)

Purpose: Selects one or more of the 19 defined test data patterns for the write-then-read error rate test.

Syntax: TP [*pattern number*] . . . [*pattern number*]

pattern The definition of the data pattern to use on each write-then-read test pass.

Description:

If more than one test data pattern is selected, a different data pattern is used on each write-then-read test pass. If you enter the TP command with no parameters, the test data patterns are displayed on the system console along with the current pattern choices.

The default set of data patterns consists of patterns 1 through 5, and 19. Patterns 1 through 7 check the controller's data collection capability. Patterns 8 through 18 test the disk media. These 10 patterns are actually rotations of two unique bit patterns designed to check each bit at every word in a sector.

You can redefine the last data pattern, number 19. This is useful if an error is suspected when a specific bit pattern is written on the disk. Refer to the discussion of the BP command for more information on setting pattern number 19.

ERTSH

When you enter a TP command with any parameters, a list of the current test patterns displays, as shown in the following example:

```
Active test patterns = 1 2 3 4 5 19
```

```
1)      0B
2)    125252B
3)    111111B    22222B    44444B
4)      66666B          155555B    133333B
5)      22222B    133333B    111111B    155555B    44444B    66666B
6)    165352B
7)    165110B
8)    165656B    42425B    127272B
9)    153535B    72727B    56565B
10)   127272B    144530B    135353B
11)   56565B    156565B    72727B
12)   153535B    127272B    165656B
13)   72727B    56565B    21213B
14)   4102B    10204B    20410B    41020B    102041B
15)   10204B    20410B    41020B    102041B    4102B
16)   20410B    41020B    102041B    4102B    10204B
17)   41020B    102041B    4102B    10204B    20410B
18)   102041B    4102B    10204B    20410B    41020B
19)   177777B
```

Execute Write/Read Error Rate Test (WR)

Purpose: Execute a write-then-read error rate test over the specified disk range.

Syntax: WR [SQ|RN]

SQ|RN The specification of sequential (sq) or random (rn) access.

Caution The WR command overwrites data on the specified LU. Be sure to back up important data before using this command. You can run a non-destructive, though less efficient, test by selecting the RO (Read Only) command instead, as follows:

```
CI> ru,ertsh,lu lu:wr:ss:ex
```

or

```
CI> ru,ertsh,lu lu:ro:ss:ex
```

Description:

The LU cannot be mounted to the file system when this test is executed. A data pattern is written on the disk and read back. The data read is then compared to the data written. The parameter defines the use of a sequential or random access, as in the RO command. The write-then-read test is more thorough than a read-only test, because more of the disk mechanism is used and worst case data can be placed on the disk.

When you enter the command in response to an interactive prompt, ERTSH asks for permission to overwrite the disk data. For example:

```
ERTSH (?? for help)> wr
```

```
All data on LU 56 will be destroyed, OK to proceed (YES or NO)? > YES
```

Using ERTSH

You can use ERTSH in several different ways to find and isolate disk errors. The following sections contain suggestions on how to use ERTSH.

Disks Formatted with FORMF

You can use ERTSH to test disk drives that were formatted using FORMF, but you cannot programmatically correct any defects that are found until you re-initialize the disk with FORMA.

Note that FORMA addresses sectors differently than FORMF. This means that you cannot use FORMA to correct defects found by running ERTSH on a disk formatted with FORMF. See the earlier section on “Converting from FORMF to FORMA” for additional information.

Initializing and Testing a New Disk

ERTSH is most effective when you use it after you use FORMA to initialize and spare a new disk. See the section “Installing a New Disk” in the earlier discussion of FORMA for the recommended procedure.

Typical Use of ERTSH

A disk problem is usually suspected when the disk driver returns an error message or when disk data appears to be corrupt. Errors usually appear when a particular disk location is being accessed. The following example illustrates a typical run of ERTSH.

While accessing a file on LU 56, a disk error is reported that you want to correct. You should back up this LU before running ERTSH, because you may need to use a write-before-read test to isolate the error. After backing up the LU, run ERTSH and specify the LU.

ERTSH

For example:

```
CI> ru,ertsh
ERTSH Rev. 2526 <850524.1705>
ERTSH (?? for help) > lu 56
```

Next, select either a write-then-read or read-only test to run. The write-then-read test is more thorough and usually detects more errors. The duration of the test and the test patterns used in a write-then-read test can be defaulted to 1, 2, 3, 4, 5, or 19. Because this test is destructive, ERTSH asks if it can overwrite the current disk data. When you enter “YES”, the program continues.

```
ERTSH (?? for help) > wr
```

```
All data on LU 56 will be destroyed, OK to proceed (YES or NO)?> YES
```

As the test executes, it reports the start of each test pass and any errors detected during that pass. The error messages include the sector in error, disk status and error register values, and an English translation of the error and status registers:

```
Write/Read Test on LU 56      Pass #1      Using test pattern #1
Write/Read Test on LU 56      Pass #2      Using test pattern #2
Disk Error on LU 56      Cylinder = 137      Head = 2      Sector = 9
Status = 121B      Error = 100B
ECC/CRC error in data field on read

Disk Error on LU 56      Cylinder = 154      Head = 4      Sector = 15
Status = 124B      Error = 0B      Data error corrected by ECC

Write/Read Test on LU 56      Pass #3      Using test pattern #3
Disk Error on LU 56      Cylinder = 137      Head = 2      Sector = 9
Status = 121B      Error = 100B
ECC/CRC error in data field on read

Write/Read Test on LU 56      Pass #4      Using test pattern #4
Write/Read Test on LU 56      Pass #5      Using test pattern #5
Write/Read Test on LU 56      Pass #6      Using test pattern #19
Write/Read Test on LU 56      Pass #7      Using test pattern #1
Write/Read Test on LU 56      Pass #8      Using test pattern #2
Disk Error on LU 56      Cylinder = 137      Head = 2      Sector = 9
Status = 121B      Error = 100B
ECC/CRC error in data field on read

Disk Error on LU 56      Cylinder = 154      Head = 4      Sector = 15
Status = 124B      Error = 0B
Data error corrected by ECC

Write/Read Test on LU 56      Pass #9      Using test pattern #3
Disk Error on LU 56      Cylinder = 137      Head = 2      Sector = 9
Status = 121B      Error = 100B
ECC/CRC error in data field on read

Write/Read Test on LU 56      Pass #10     Using test pattern #4
Write/Read Test on LU 56      Pass #11     Using test pattern #5
Write/Read Test on LU 56      Pass #12     Using test pattern #19

Cleaning up disk test area
```

After completing 12 test passes (the default pass count) and cleaning up the test disk area, ERTSH reports a summary of the errors:

```
Error Summary for LU 28
```

```
Table entry format      Most severe error code - Error count for track
```

- 1 Soft error: ECC corrected error
- 2 Hard error: ECC/CRC error, ID not found, etc.
- 3 Data miscompare undetected by controller
- 4 Error on EXEC call to the disk

```
Cleaning up disk test area
```

Cyl./Sc	Head Number					
	0	1	2	3	4	5
137/ 9			2 4			
154/15					1 2	

An error summary entry contains a general error code (1-4) followed by the total number of times an error was detected at that sector. In the example above, the error summary indicates that there are two defective sectors. The first error (cylinder 137, head 2, sector 9) has an error code 2 or a hard error that occurred 4 times. The error is so severe that the data is probably not recoverable and should therefore be spared.

The second entry (cylinder 154, head 4, sector 15) indicates the error correction code (ECC) is correcting the error, so the data read from that sector may still be valid. Even though data can be written and read correctly (using error correction), you should spare the sector before the defect gets worse and becomes a hard error. For more information, refer to the “Error Interpretation” section, later in this chapter.

You can spare the defective sectors by using the SS command from ERTSH, or you can exit the utility and use FORMA directly. In the following sequence, the SS command is used to have ERTSH schedule FORMA to spare the defective sectors. Since the default SS command parameter is used, all the defective sectors are spared.

```
ERTSH (?? for help) > ss
```

```
Sparing sector at Cylinder = 137      Head = 2      Sector = 9
FORMA: INTEGRATED DISK FORMATTER, REV 2526 <850530.0911>
FORMA: SPARING COMPLETE
```

```
Sparing sector at Cylinder = 154      Head = 4      Sector = 15
FORMA: INTEGRATED DISK FORMATTER, REV 2526 <850530.0911>
FORMA: SPARING COMPLETE
```

```
ERTSH (?? for help) > ex
```

Note that it takes about two hours for this sequence to test for errors, isolate them, and spare the defective sectors. When the test completes, you can exit ERTSH and restore the files from your backup device.

ERTSH

If ERTSH detects an error in a command while processing a string of commands, it displays an error message and ignores the rest of the commands. In interactive mode, ERTSH returns with another prompt or terminates if it is processing runstring commands.

Break Detection

You can enter the RTE BR (Break) command to terminate an error rate test before the defined number of test passes complete. Upon detecting the break, ERTSH cleans up the disk and reports the error summary to the terminal. Setting the test pass counter to zero (infinity) and running a test that you terminate with the BR command is an easy way to perform a lengthy test.

Specialized Use of ERTSH

If a disk error is encountered while a program is attempting to read a certain track from a disk LU, you can limit the disk test area to reduce the test time because you know the general disk location.

First, translate the track address to an absolute cylinder and head address. After selecting an LU with the LU command, use the CY and HD commands to display the physical cylinder and head bounds for the given LU. You can use the following calculations to get a precise disk area to test, or you may be able to approximate the location of the defect and select an appropriate disk test area, as shown in the following example:

```
1st LU track = (1st LU cylinder) * (# of heads on disk)
absolute track = (1st LU track) + (track # within track)
absolute cylinder = (absolute track) Div (# of heads on test LU)
absolute head = (absolute track) Modulo (# of heads on test LU)
```

Some of the above values can be found in the system generation listing. Since some defects cross track boundaries, you should also test the tracks adjacent to the suspected track.

For example, if the suspected disk location is somewhere on the track at cylinder 23, head 2, it is appropriate to test the area from cylinder 20 to cylinder 25. You can test all the heads or just head 2, depending upon how thorough you want the test to be. Testing a small portion of the disk is much faster than testing the entire disk.

In deciding how thorough a test to perform, consider test duration and the selected test pattern (if the test is a write-then-read test). As shown in the sample listing, a disk defect may not be detected on every access to that sector. In general, the smaller the defect and the closer it is to the start or end of a sector, the less likely it is to be detected on each pass of ERTSH.

Another factor related to detecting an error is the bit pattern written on the disk. Some defects appear to be pattern sensitive and may require using several different patterns before the defect is isolated. ERTSH contains 18 standard test patterns and one user-definable test pattern number (19).

For the write-then-read test, a track size buffer is created by duplicating a sector buffer built by repeating the values in a given pattern. For example, test pattern 3 consists of the three values 111111, 222222, and 444444 octal.

The sector buffer built from these values is as follows:

```
111111B 22222B 44444B 111111B 22222B 44444B 111111B
```

and so on. This particular pattern corresponds to a bit pattern of a binary 1 followed by two binary zeros (100100100100...).

It can be very difficult to select a test pattern. There is no way to know ahead of time which pattern will detect a given error. However, it is known that severe errors are usually detected by several test patterns. The test pattern defaults usually detect defects, but you may need to use other patterns if a particular defect is not isolated.

Patterns 6 and 7 stress the controller's ability to distinguish between data bits within a word. Patterns 8-13 and 14-18 are permutations of two basic test patterns that are rotated to check all bit positions. An ideal test would run several passes with each test pattern, but this becomes prohibitive due to the time required.

You can use the user-definable pattern (number 19) if you suspect that data has been corrupted and the disk defect is pattern sensitive. For example, a file is supposed to contain the value 100 in a given word, but the value 101 is returned when the file is read. A test pattern of 100 causes the value of 100 to be written, then read and compared on the whole disk range. The pattern can be defined as follows:

```
ERTSH (?? for help) > bp 100
```

Installation Procedure Summary

Here is a summary of steps to take when you install a new internal hard disk or run FORMA/ERTSH (except for Step 1) on a disk currently in FORMF format.

1. Find the label attached to the top of the disk drive that lists the areas that the manufacturer has found to be defective. Record the *cyl*, *hd*, and *byte/index* information. This data is used in Step 3, in the FORMA runstring.

Caution The initializing process overwrites data on the entire disk. Be sure to back up your data before initializing disks currently in use.

New disks are shipped from HP already initialized and spared in FORMA format, and need not be initialized again. If you are using new disks, skip to Step 4.

2. Initialize the hard disk by running FORMA, using the IM command. If the drive was never formatted (or is currently in FORMF format), use the N option to prevent spares from being retained, to save time. If the drive was already formatted and spared using FORMA, use the P option to save the old spares.

```
CI> ru,forma,im,<lu><n|p>
```

ERTSH

This should take about 40 minutes for a 15 MByte drive.

3. Spare the defects by using the FORMA SS command and entering the data recorded in Step 1.

```
CI> ru,forma,ss,<lu>,n,b,<cyl>,<hd><byte/index>
```

4. Test for new defects by running ERTSH. Using the SS command, ERTSH locates the defects and programmatically calls FORMA to spare them.

Error Interpretation

ERTSH reports all disk errors (including those that cannot be corrected). Each error is reported when it occurs during an error rate test, and a summary of error locations is listed at the end of each test.

The error messages contain the cylinder, head, and sector where the error was detected along with the status and error code. For example, if an ECC error is detected, the following error is displayed:

```
Disk Error on logical LU 16   Cylinder = 113   Head = 1   Sector = 3
Status = 124B   Error = 0B
Data error corrected by ECC
```

The first line indicates the location of the error. This information is needed if the defective disk area is to be spared later. The second line contains the disk status and error register returned from the controller. Note that this is not the same status that is returned from the driver on a standard request to the disk driver or EXEC call. The last item in the error message is an English translation of the error status and error bytes.

The other type of error message indicates a data miscompare. Again, the disk location is identified, along with the data which caused the miscompare.

```
Disk Error on LU 16   Cylinder = 101   Head = 2   Sector = 13
Miscompare word = 11
Value written = 0000000000000000   0B
Value read    = 0000000000000100   4B
```

The word in the sector that caused the error is identified along with the expected and actual value read from the disk. The values are first given as binary, then octal values.

Occasionally, ERTSH reports an error with a sector number of “?”. This happens when ERTSH detects an error on the track but cannot determine the sector that is at fault. If “?” also appears in the error summary, check the number of times the error occurred. If the total number of errors for that sector is only one or two, the error was probably transient and you can ignore it. If the error count is large, retest the area around the track in question.

Note that ERTSH reports error conditions that the disk controller or drive may mask from you. For example, even if the disks ECC/CRC correct an error in disk data, and the disk data is now correct, ERTSH still reports an error condition.

At the end of a test, an error summary is created to help in error interpretation. Errors are grouped into four categories when placed in the error summary, as shown in the following sample output.

```
Error Summary for LU 56
```

```
Table entry format  Most severe error code - Error count for track
```

- 1 Soft error: ECC corrected error
- 2 Hard error: ECC/CRC error, ID not found, etc.
- 3 Data miscompare undetected by controller
- 4 Error on EXEC call to the disk

Soft errors are data reads from the disk that were corrected by the disk controller. This ECC correction is a standard mechanism used in many disk drives to compensate for minute defects in the disk media. Even though the data in these sectors is not corrupt, you should spare the sectors before a defect becomes a hard error.

Hard errors are more severe than soft errors, and are usually easier to detect. Hard errors can be caused by a disk defect in the sector data area or sector header area or by a defective HP 12022A interface card. The following errors are all classified as hard errors:

```
Bad block mark found in ID field
ECC/CRC error in data field on read
ID field for given location not found
Unable to seek to track zero
Data address mark not found after ID field
Last command aborted due to error
```

```
Card is still busy
Drive not ready
Write fault
Seek not complete
Disk is write-protected
```

In the messages above, the first six errors are encoded in the error register value, and the last five are encoded in the status register.

Data miscompare errors indicate that a malfunction of the disk or disk controller was not detected by the controller but was detected by ERTSH. This error usually indicates a defective HP 12022A controller.

An error on EXEC call may be caused by a defective HP 12022A interface or a defective disk; or the problem can be simpler, such as not having the disk cables on tightly or not having the disk powered up.

ERTSH

When you interpret the table as a whole, you must consider other factors. For example, an error table like the following one may indicate one of two likely errors.

Cyl. / Sc	Head Number							
	0	1	2	3	4	5	6	7
104/ 0					4	4	4	4
104/ 1					4	4	4	4
104/ 2					4	4	4	4
104/ 3					4	4	4	4
104/ 4					4	4	4	4
104/ 5					4	4	4	4
104/ 6					4	4	4	4
104/ 7					4	4	4	4

The error summary above indicates that every sector on every track using head 4 and 5 is defective. The real problem may be that heads 4 and 5 of the disk drive do not work, or they may not even exist. A 10 MByte disk (which has 4 heads) may have been used instead of a 15 MByte disk (which has 6 heads).

In this case, it does not make sense to spare every defect shown in the table, because the real problem is not defective sectors. Whenever you see a large number of errors in the summary, look for causes other than defective sectors.

Error Messages

When ERTSH detects an error, it first prints:

```
ERTSH Error: Error in xx command entered
```

This indicates which command in the sequence caused the error and is followed by one of the messages listed below.

ERTSH Error: Command requires superuser capabilities

Commands that cause disk data to be destroyed can only be executed by users who have superuser capability.

ERTSH Error: Defects on microfloppy disk cannot be spared

Defects on microfloppy drives cannot be spared using ERTSH.

ERTSH Error: Disk LU xx is mounted to file system Command can not be executed.

Commands that cause disk data to be destroyed cannot be executed while the disk LU is mounted to the file system.

ERTSH Error: Error in call to LU xx Error = xxxx

An error was detected during an I/O request to the given LU. The error code xxxx is a system error code such as IOTO (device timeout).

ERTSH Error: Error summary table is full, error not in summary

The error summary table is full. Errors are still reported, but may not appear in the error summary at the end of the test pass.

ERTSH Error: Illegal Disk LU

The LU you entered is not an integrated disk LU.

ERTSH Error: Internal Error xx

ERTSH encountered an unexpected situation. Please report this error to your local HP Representative.

ERTSH Error: Invalid parameter

The parameter you entered is inappropriate for the command.

ERTSH Error: IO error on device LU xx

There was an error on the I/O request to the given device.

ERTSH Error: Lower bound greater than upper bound

You specified a range in which the lower bound was larger than the upper.

ERTSH

ERTSH Error: No test LU selected for test

You must select a test disk LU before a test pass can be executed.

ERTSH Error: No tracks on test LU lie within cylinder/head range

The head and cylinder range do not contain any tracks from the test LU. Redefine the cylinder and/or head test range.

ERTSH Error: Parameter *xx* is not a valid integer

The parameter you entered is not a valid integer.

ERTSH Error: Unable to lock Disk LU *xx*

ERTSH must be able to lock the test LU in order to perform the test.

ERTSH Error: Unable to unlock Disk LU *xx*

ERTSH was unable to unlock an LU it had previously locked.

ERTSH Error: Unknown command, type ?? for help

You specified an unknown or invalid command mnemonic.

ERTSH Error: *xx* is out of range

The parameter you entered is too large or too small.

FC

FC allows you to copy FMGR files between disk cartridges and tape media, either disk-to-disk, disk-to-tape, or tape-to-disk. Tape-to-tape copying is not supported. FC does not support LUs greater than 63. You can copy the files to/from a 1600-bpi magnetic tape or a CS/80 cartridge tape drive.

You may give a file being copied a different file name, security code, or cartridge by specifying those fields of the destination name. When you copy from disk, file extents are automatically gathered and copied in ascending order following the main extent. Optionally, you may eliminate all extents and copy all sections of a file to the main extent.

You may also select command options that purge source disk files after copying, list or suppress the listing of files copied, replace duplicate files with the last duplicate copied, or verify the copy. (The Verify option is selected in addition to the tape checksums, which are always used.)

Calling FC

You may run FC from a terminal or from a command file. Any single FC command can be entered in the runstring. When you run FC from a command file, the FC transfer (TR) command lets you specify a command file as the FC utility command source.

Interactive mode is entered by invoking FC with no runstring. Commands are entered at the FC: prompt. FC executes each command before prompting for the next one, continuing until you enter the EXIT command. For example:

```
CI> ru,fc
FC: command
FC: command
.
.
.
FC: ex
:
```

You may enter just the first two characters, or up to the whole name of a command; for example, EX, EXI, and EXIT all cause FC to exit. Uppercase and lowercase are acceptable. You may enter a comment at the end of a command string, or by entering an asterisk (*) as the first character in the response. If you enter a comment at the end of a command string and do not use all the available optional parameters for the command, you must insert commas at the end of the string as placeholders for the omitted parameters.

When a command is entered in the runstring, as:

```
CI> ru,fc,command
```

FC executes the command and terminates.

FC

To run FC from a control file, enter the TR command in the runstring:

```
CI> ru,fc,tr control_file
```

The *control_file* parameter can be a device LU or a file name. FC executes each of the commands contained in the control file, then terminates.

FC Commands

FC commands let you configure the copy operation and specify names for the formatted tape files and comment files, selectively copy files and group copy commands, transfer to and return from command files, and direct listings to a log or list device.

Configuration commands set attributes that affect the execution of other FC commands. They let you select the list device, enable/disable command echoing to the list device, set the title and comment file for tapes subsequently written, and set the scratch file disk cartridge.

The Copy (CO) command initiates the copy operation from disk-to-tape, tape to disk, or disk-to-disk. Cartridges and tapes can be written and restored by naming the devices as the source and destination parameters. Files being copied can be selected by name, security code, cartridge, and so on, as desired. Selection by name can make use of wildcard characters.

In copying to or from tapes, there may be applications that require specifying multiple source/destination parameters in a single copy operation. This is true, for example, if you are copying files to a tape and renaming them, or if tape files are to be copied to disks other than those specified in the tape directory. Grouping copy commands lets you combine multiple copy commands into a single copy operation.

Listing commands direct listings of the FMP cartridge list; global FMP cartridge list; tape directory list; and the tape comment and header files to your terminal, a selected list device, or a file.

Transfer and exit commands let you transfer control to and from a control file and stop FC operations.

Table 2-14 summarizes the commands, which are described in the following sections.

Table 2-14. FC Commands Summary

Commands	Description
Information Command	
?	Provides a summary of commands and command syntaxes
Configuration Commands	
CF Name Comment File <i>namr</i>	Establishes <i>namr</i> for tape comment file
ECho ON OFF	Turns echoing of commands to list device ON or OFF
LL Set list device <i>namr</i>	Sets list device or file (destination for listed messages and commands)
SC cratch Files <i>cart</i>	Specifies disk cartridge to be used for internal scratch files
Title <i>title</i>	Establishes title to be used in tape header file
Copy and Related Commands	
CO py <i>srce</i>	Copies files as specified by parameters
DE fault <i>srce</i>	Sets default source, destination, and options for subsequent COPY commands
GR oup CO Commands	Groups more than one copy command
AG Abort Group	Terminates grouped copy commands execution
EG End Group	Causes execution of the copy commands between the preceding GR command and this EG command
Listing Commands	
CL Cartridge List <i>-thu</i>	Lists FMP cartridge list, or cartridges included on tape
DL Directory List <i>srce</i>	Compiles directory list of FC tape
LC List Comment File <i>-thu</i>	Lists comment file from FC tape
LH List Header File <i>-thu</i>	Lists header file from FC tape
Transfer and Exit Commands	
AB ort	Aborts FC, including any active group copy
EX it	Exits FC (if a group copy is active, it is processed before FC is aborted)
TR ansfer <i>file</i>	Transfers to/from FC command file
Comment Command	
*	Identifies following string as comment line

FC

Command Summary Function (?)

Purpose: Writes the FC command summary to the terminal.

Syntax: ?

Description:

When you enter ? at the prompt, the FC command summary shown below is written to the terminal.

If you enter ? together with another command or with an option (for example, "? , SCRATCH"), FC displays additional information about the specified command or option.

```
----- FC commands -----  commands may be abbreviated to 2 chars

COPY, srce, dest, opts, [file1], [file2], [msc] copy files
DEFAULT , srce , dest , opts          set defaults for COPY command
GROUP / EG / AG                       begin / end / abort GROUP of COPY commands
LL , namr                              set list file/device (dash means log device)
DL , srce , [msc] , opts              list tape directory (srce = -tlu or -tlu{namr})
CL , [-tlu] , options                 list local cartridge list or tape cartridge list
CLAL                                  list global cartridge list (RTE-IVB/6 only)
LH , -tlu , opts                      list tape header file
LC , -tlu , opts                      list tape comment file
ECHO [, ON/OFF]                       turn ON/OFF cmd echo to list device (default ON)
TITLE , title                         set tape title (for subsequent COPYs to tape)
CF , comment-file-namr                set comment file (for subsequent COPYs to tape)
TR, namr                              transfer to cmdfile/device(dash means log device)
TR                                    return from command file/device
EXIT                                  exit FC
ABORT                                 abort FC (same as EX except if copy group active)
SCRATCH , cartridge                   set cartridge that FC will use for scratch files
* comment                             command line starting with * treated as comment
  HELP [, key [, lu                   get help, RTE-IVB/6 only (useful for FMGR errors)
? , <command>                         list info about particular command(incl. options)
? , <option>                           list info about particular option (all commands)
```

Figure 2-1. Example of FC Command Summary Listing

Abort (AB)

Purpose: Stops FC execution immediately.

Syntax: AB

Name Comment File (CF)

Purpose: To specify the name of an optional comment file to be copied from disk to tapes created by subsequent copy commands.

Syntax: CF, *namr*

namr The file name of the comment file.

Description:

When the copy is to magnetic tape, the file is copied uncoded. Because CTD tapes are formatted in fixed length blocks, files must be coded for copy to CTD devices.

Comment file records should not end in a backslash. Any comment file records longer than 128 words are truncated. Note that a comment file should not contain zero length records, or a checksum error may result when you list the file with the LC command.

Since the comment file is not encoded on magnetic tapes, you can read the file using the FMGR ST, LI, or DU commands.

Cartridge List (CL)

Purpose: Displays the list of all cartridges you can access.

Syntax: CL, *-thu* [, K]

-thu Lists the cartridges stored on the tape. For example:

```
title: nnnnn
volume: x
date and time of creation: hh:mm ddd., mmmm, yyyy
created under account name: nnnnnn.nnnn
```

```
LU CRN LABEL P/G/S USER.GROUP
xx xxx xxxxxx x xxxx.xxxx
```

K Keeps the tape online.

Description:

If the cartridge was renamed as part of the copy operation, only the cartridge CRN is listed.

You can specify the K option to keep the tape drive online under all circumstances. Note that while the tape is online, it may be overwritten by another program, since it is unlocked.

You may use the LL command to direct the CL listings to another list device.

<i>file1</i> <i>file2</i>	<p>The optional limits on the range of files to be copied from the source. Within the range specified, files that conform to the source parameter are selected for copying to the destination. Note that disk resident files are referenced by namr, and tape resident files are referenced by disk file reference number. (Refer to the DL command.) If only <i>file1</i> is specified, the range is interpreted as falling between <i>file1</i> and the last file on the source medium. If only <i>file2</i> is specified, the range is interpreted as all files between the first file on the source medium and <i>file2</i>.</p> <p>If <i>file1</i> is specified but not found, no files are copied. If <i>file2</i> is specified but not found, all conforming files between the range of <i>file1</i> and the end of the source medium are selected for copying to the destination.</p> <p>If you specify multiple disk cartridges in the source parameter, the <i>file1</i>, <i>file2</i> parameters are applied separately to each cartridge; for tape sources, the <i>file1</i> and <i>file2</i> parameters apply to the entire tape or set of tapes.</p>
<i>msc</i>	<p>The system master security code. Msc enables you to specify a security code in the destination even if you do not specify a security code in the source parameters. This parameter also allows you to purge files using the P and D options even if the file security codes do not match. The msc is also required if you specify the C or ! options. (Refer to the section in this chapter on “Copy Command Options” for a detailed description of the available options.)</p>

Description:

You can write and restore cartridges and tapes by naming the devices as the source and destination parameters. Select files to be copied by name, security code, cartridge, and so on, as desired. Selection by name can make use of wildcard characters.

FC

CO Command Source and Destination Parameter Considerations

When you form the source and destination parameters, consider the following characteristics of FC:

- A null security code field in the source parameter matches files regardless of the security code. A zero in the security field matches only files with a security code of zero (unprotected files). Note that in FMGR, unprotected files are created by specifying a zero security code or by omitting the security code from the namr.
- For disk sources, a null cartridge field in the source namr is interpreted to mean the first cartridge that contains a file that matches the source file namr (including wildcard characters). You cannot use a given namr to select files from more than one cartridge.
- For tape sources, a null cartridge field in the source namr is interpreted to mean all cartridges on the tape.
- All files selected by the source and the *file1* and *file2* parameters are copied, except for type 0 files and files with “illegal names”. Files with illegal names include scratch files with numeric names.
- In the destination parameter, you may specify a file security code field only if you specify the security code in the source namr, or if you specify the master security code.
- When the specified destination cartridge differs from the one in the source, the copy operation is as follows:
 - Disk-to-disk: The file is copied to the cartridge named in the destination parameter.
 - Disk-to-tape: The file is labeled on the tape with the cartridge named in the destination parameter.
 - Tape-to-disk: The file is copied to the cartridge named in the destination parameter.
- When you copy to tape, if the destination cartridge that will be listed in the tape directory is specified by the destination parameter, it must be specified as CRN, not –LU.

CO Command Options

You may specify any of the following options with the CO command, in any order. Spaces in the option string are ignored. Table 2-15 summarizes the CO command options, followed by a discussion of each option.

Table 2-15. FC CO Command Options Summary

Options	Description
Brief, Full Status	Displays status of namrs as each file is copied, when an error occurs (B) or for successfully copied files as well (F)
Clear Destination Disk	Clears destination disk before the copy begins
Duplicate Files	Purges original file and replaces it with the duplicate
Eliminate Extents	Combines all extents into the main extent
Ignore Data Errors	Overrides checksum or verify errors and allows the file to be copied
Keep Tape online	Keeps tape unit online in all circumstances
Lock, Open	Locks a cartridge or opens individual files for the copy
Purge Source File	Purges a source file after it is copied
Single Volume Copy	Allows you to mount a single tape volume set
Tape Length Display	Displays required tape length for copying specified files
Unused Space	Allows FC to use more space on the destination medium than required for files being copied, to speed up copying
Verifying Data	Reads data and compares checksums to verify that data is correct

Brief, Full Status Display Format (B,F)

You can have the source namr (and sometimes the destination namr) displayed as each file is copied, if desired. If you select Brief, the namrs are displayed only when an error occurs, to identify the affected file. When you select Full, the namrs are displayed for successfully copied files as well as those with errors. (Refer to the “Error Messages” section later in this chapter for the format of the displayed namrs.) Brief mode is the default when you select files only by tape LU and/or disk cartridge; Full mode is the default in all other cases.

FC

Clear Destination Disk (C,!)

When you select this option, the destination disk is cleared before the copy begins. Note that the boot extension file BOOTEX is not cleared if it is the first entry in the disk directory. The Clear option is permitted only if the destination is a disk cartridge and the system master security code (msc) is correctly specified in the FC command string. During the session, any value in the msc parameter is accepted for clearing private and group cartridges. The true system master security code is required only for system and non-session cartridges.

Before clearing the disk, FC issues the message:

```
Do you really want to purge disk LU nn, CRN nnn ?
```

and waits for a Yes or No response.

The ! specification works the same as the C option, except that the above message is suppressed.

Replace Duplicate Files (D)

Normally, when a duplicate file is encountered in a copy, a duplicate name error (FMGR-002) is issued and the duplicate file is not copied. However, if you select the Duplicate option, when a duplicate file is encountered, the original file is purged and replaced with the duplicate.

Typically the replacement does not take place unless the security code of the duplicate destination file matches that of the destination file to be replaced (according to the standard FMP definition of matching security codes). However, if the system master security code is correctly specified in a CO command that includes the D option, file security code rules are superseded, and all duplicates are replaced regardless of whether the file security codes match. (The security code of the new destination file comes from the source file except when this is overridden by a security code specified in the destination parameter.)

Duplicate names can occur when more than one source cartridge is copied to a single destination cartridge. In such cases, the file from the first source cartridge normally is copied successfully, and the rest are not copied because of duplicate name errors. If the D option is selected, each successive duplicate replaces the previous one, so that only the last duplicate file is copied.

Eliminate Extents (E)

This option combines all file extents into the main extent. Unused blocks are not truncated. Extents are not eliminated in files that have missing extents (“sparse” files), or if the resulting main extent would exceed 16383 blocks. In these cases, a warning is issued. If subsequent changes to a file result in the creation of a new extent, the extent created is the same block size as the main extent.

Ignore Data Errors (I)

In general, a file is not copied if a checksum or verify error occurs. The I option overrides this feature on tape-to-disk copies, and the file (including data for which checksum errors were detected) is copied to the destination. Files that are copied with errors are not identified as such in the directory entry for the file; however, a message is issued that defines the range of bad blocks copied. The I option is applicable only for tape-to-disk copy operations.

Keep Tape Online (K)

This option keeps the tape unit online in all circumstances. While this option is a convenience, it makes it possible for another program to overwrite the tape because the tape unit remains unlocked and online. Refer to the “Tape Handling” section in this chapter for more details.

Cartridge Lock, Open (L,O)

You can copy files to/from cartridges either by locking the cartridge (Lock mode) or by opening individual files for the copy (Open mode). Lock mode is generally faster, but excludes other programs from accessing the cartridge. Refer to the “Performance Considerations” section in this chapter for details. Lock mode cannot be used if other programs have files open on the cartridge, and should not be used if other programs may need access to files on the cartridge during the copy. Once a Lock mode copy is started, other programs cannot open files on the cartridge for the duration of the copy.

In Open mode, source files are opened non-exclusively, and destination files are opened exclusively. Therefore, even the Open mode restricts access to files, although the restrictions are less severe than those in Lock mode. If neither L or O mode is specified, the defaults are as follows:

- If the source name (or names) contains wildcards or is null, or if more than one copy operation is grouped, a cartridge lock is attempted. If it is successful, Lock mode is used.
- If the lock is rejected, or if there are no wildcard or null source names and the GR command is not used, Open mode is selected.

Purge Source File (P)

When you select the P option, FC purges the source file after it is successfully copied to the destination. To purge a file with a non-zero security code, you must correctly specify either the security code field of the source parameter, or the msc parameter; otherwise, the file is not purged. This does not, however, prevent the file from being copied.

If the destination is a tape, the source files are purged after all files are copied. Thus, in the event of an error, or if you abort the disk-to-tape copy while it is in process (using the break command), the source files are not purged.

Copy Single Volume of Multi-Volume Tape Set (S)

Usually, when you copy from a multiple volume tape set, all volumes must be read in sequence. When you select the Single option, this feature is suppressed, and a single tape volume set can be mounted. Note that you cannot copy files that cross volumes when you use the S option.

Display Required Tape Length (T)

When you select this option, FC reads the disk directory, calculates the amount of tape required to copy the specified files, and issues the appropriate message.

FC

Recover Unused Space (U)

In order to copy data faster, FC may use more space on the destination medium than is required for the files being copied. This can only happen on disk-to-tape copies and lock mode disk-to-disk copies. Although this does not typically result in a significant waste of space, specifying the U option prevents any extra space from being used.

Verify Transferred Data Integrity (V)

The V option specifies the following data verification operations:

Disk Read	Data is read twice and the checksums are compared to verify the read.
Disk Write	Data is written to the disk and then immediately read to a second buffer. The two data buffers are compared to verify the write.
Tape Read	The checksum is calculated for each buffer read from the tape and compared to the checksum read with the data. (This tape read verification is performed even if the V option is not selected.)
Tape Write	The tape is rewound after each volume is written, and a tape read verification performed as above.

CO Command Examples

The following examples show how to use CO to copy files between disks and tapes.

Example 1: Copy cartridge 10 to tape LU 8.

```
co,10,-8
```

Example 2: Copy all files beginning with & on cartridge 10 to tape LU 8.

```
co,&-----::10,-8
```

The ----- characters in the name field define the wildcard characters. The null security code matches any file security code.

Example 3: Copy three files to cartridge 20.

```
co,{filea,fileb,filec},20
```

Example 4: Copy file1 from cartridge 30 and file2 from cartridge 40 to tape LU 8, and verify the data transfer.

```
co,{file1::30,file2::40},-8,v
```

Any files existing on LU 8 are overwritten as the specified files are copied.

Example 5: Copy all type 4 and type 5 files from cartridge 10 to tape LU 8.

```
co, {::10:4, ::10:5}, -8
```

Example 6: Copy all files that begin with % and & from cartridge 10 to tape LU 8.

```
co, {%------:10, &-----:10}, -8
```

Example 7: Restore all files from tape LU 8 to their original cartridges.

```
co, -8
```

Since the destination parameter is omitted, the source namrs (from the tape) are used for the destination files. The tape directory file identifies the cartridge from which each file was copied to the tape.

Example 8: Restore all files from tape LU 8 to cartridge 10.

```
co, -8, 10
```

The tape directory information identifying the cartridge of origin is ignored, and all files are copied to cartridge 10.

Example 9: Restore from tape LU 8 only those files that were copied to the tape from CRN 10.

```
co, -8 {::10}
```

Because no destination cartridge is specified in the command, the files are restored to the same CRN.

Example 10: Restore from tape those files that were copied to tape from CRNs 10 and 20.

```
co, -8 {::10, ::20}
```

Because no destination cartridge is specified in the command, the files are restored to their original CRNs.

Example 11: Copy file1 to file2 on the same cartridge.

```
co, file1, file2
```

FC

Example 12: Copy file A on cartridge 10 to cartridge 20, and name the file “B”.

```
co,A::10,B::20
```

Example 13: Copy file A to the same cartridge, name it file B, and change the security code from 2 to 3.

```
co,A:2,B:3
```

Example 14: Copy all files within the range of file2 and file6 from cartridge 10 to cartridge 20.

```
co,::10,::20,,file2,file6
```

If file2 is not found, no files are copied. If file6 is not found, all files in the range of file2 to the end of the source medium are copied. In specifying the file1,file2 parameters, the files must be specified in ascending order of occurrence on the medium. In other words, no files are copied if file2 precedes file6 on the cartridge when you issue the following command:

```
co,-8{fileA},::20,,file6,file2
```

Example 15: Copy all files beginning with % within the range of fileA and %file6 from cartridge 10 to cartridge 20.

```
co,%-----::10,20,,fileA,%file6
```

Example 16: This example is the same as the one above, except that the source is a tape -LU and, therefore, the range of files must be specified as file reference numbers obtained from the DL command.

```
co,-8{%-----},20,,2,6
```

Example 17: Copy all files from file 1 through file 5 from tape LU 8 to cartridge 10.

```
co,-8,::10,,,5
```

Default (DE)

Purpose: Used with the CO command to set default values for the source, destination, and/or options parameters.

Syntax: DE [, *srce* [, *dest* [, *options*]]]

srce, dest, options Defined for the CO command, except that a list of namrs may not be used in the *srce* parameter.

Description:

If you do not enter source and destination parameters in the CO command, the values you specified for those fields with the DE command are used. You can override the DE command values by specifying the source and/or destination parameters.

When you include options in the DE command, they are appended to any options specified in succeeding CO commands. For example:

```
de , , , epv
co , srce , dest , cde
co , srce , dest , d
```

In this command, C, D, E, P, and V are specified for the first CO command, and D, E, P, and V are for the second CO command. If you specify conflicting options with DE and CO, the CO option overrides the one selected with DE.

Each DE command supersedes the preceding DE specification; therefore, you can cancel defaults by entering a DE command with all null fields.

FC

Directory List (DL)

Purpose: DL provides a directory list of files on an FC tape specified by the command. (For directory listings of disk cartridges, use the FMGR DL command.)

Syntax: DL, *srce* [, *msc* [, *options*]]

srce The source FC tape device, with optional *namr*, that specifies the files to be included in the list. If you omit the *namr*, all files are listed. The source can only be a tape; use the FMGR DL command for listing cartridge files. If you specify a *namr* with the negative tape LU, you must enclose the optional *namr* within braces, as:

-tlu {*namr*}

msc The system master security code. When the *msc* is specified, the file security code is contained in the directory list.

options One or more of the following options:

- F Selects Full option, which results in including the extent size, record size, and # extents columns in the directory list.
- S Selects Single option, which compiles a directory of only one volume of a multi-volume tape set.
- K Keeps the tape unit online under all circumstances. Note that while the tape is online, it can be overwritten by another program, since the online tape is unlocked.

Description:

You can direct the DL listing to another list device by using the LL command. The tape directory list format is shown in the following example.

```

title: SAMPLE OF TAPE DIRECTORY FORMAT
volume: 2  date and time of creation: 7:09 PM TUE., 14 JULY, 1992
created under account name: KURT.SYSTEM

name      crn      LU      type      extent  record  #extents  security  diskfile
          size    size    log      phy      code     ref #

JLMHI    100      25      1         1        128      1         1         HP         1
*         100      25      4         1         1        1         1         HP         2
DVA32    100      25      4         11        1        1         1         HP         3
*****
mounted volume starts here
*****
'DVM33    100      25      4         92        3         3         HP         4 CONT
*****
mounted volume ends here

*****
'DVM33    100      25      4         92        3         3         HP         4 CONT
'CM80L    100      25      4         50        1         1         HP         5
**         100      25      4         1         1         1         HP         6
'VERIF    100      25      4         24        1         1         HP         7
'PBERS    100      25      4        234        1         1         HP         8
'BKUP     100      25      4         55        1         1         HP         9
&BOOTC    100      25      4         18        1         1         HP        10
/BIGL     100      25      4          2         1         1         HP        11
'DISK     100      25      4          7         1         1         HP        12
'R        100      25      4          3         1         1         HP        13
'NUMS     100      25      4         41        1         1         HP        14
-----

```

NOTES:

1. title – specified using TI command
2. crn, LU – appear only if multiple cartridges are on a tape
3. extent size, record size – appear only if F option specified
4. # extents – appears only if F option specified:
 - log = logical # of extents = last extent + 1
 - phy = physical # of extents = total extents allocated to file
5. security code – appears only if msc specified correctly
6. “mounted volume starts/ends here” messages – appear only in multi-volume directory lists
7. CONT – appears if file is continued on next volume or is continued from previous volume

Figure 2-2. Example of Tape Directory List Format

FC

Echo Command (EC)

Purpose: To have commands echoed to the list device as the command is processed, or to suppress echoing.

Syntax: EC , ON | OFF

Description:

Initially, command echoing to the list device is OFF. The EC command specifies that each command is to be echoed to the list device as the command is processed and, subsequently, to suppress echoing if desired. The default is ON if neither ON or OFF is specified.

Exit (EX)

Purpose: Exits FC. If a group copy is active, it is processed before FC terminates.

Syntax: EX

Group Copy Commands (GR, EG, and AG)

Purpose: To combine multiple copy commands into a single operation.

Syntax: gr
co , *parameters*
co , *parameters*
.
.
.
co , *parameters*
eg

Description:

When FC encounters the GR command, it keeps track of all subsequent CO commands but does not execute them immediately. Instead, FC executes them as a single operation when it encounters the EG (End Group) command. You can abort the GR operation by using the AG (Abort Group) command, which causes the group to be terminated before any of the CO commands are executed.

If the grouped CO commands have a common source or destination, you may use the DE command to advantage. See the following example.

DE , -8	(Common source tape -LU)
gr	
co , ::10 , 40	(Restore CRN 10 from tape to CRN 40)
co , ::20 , 50	(Restore CRN 20 from tape to CRN 50)
co , ::30 , 30	(Restore CRN 30 from tape to itself)
eg	(End group)
DE	(Cancel default)

DE, , -8	(Common destination tape -8)
gr	
co, abcd, wxyz	(Copy file ABCD to -8 as WXYZ)
co,efghij, qrstu	(Copy file EFGHIJ to -8 as QRSTU)
EG	(End group)

List Comment, List Header Files (LC, LH)

Purpose: The LC and LH commands let you list the comment file (LC) and the ASCII header file (LH) to the list device.

Syntax: LC, *-thu* [,K]

and

LH, *-thu* [,K]

Description:

You may use the LL command to direct the file listing to another list device.

You can specify the K option to keep the tape drive online under all circumstances. Note that while the tape is online, it can be overwritten by another program, since the online tape is unlocked.

List Device (LL)

Purpose: Lets you specify a device or file other than the terminal to receive list information.

Syntax: LL, *namr* | -

namr The *namr* of the device or file to receive the list information.

- Resets the list device to the terminal; used if a preceding LL command named another list device.

Description:

Initially, the list device for FC is the terminal. If you specify a file that does not already exist, it is created. Any data in an existing file is overwritten with the new information.

You may use the LL command to list cartridge and tape directories, header files, comment files, “?” command responses, and command echoes. Copy status messages and error messages are always displayed on the terminal.

FC

Scratch Area Definition (SC)

Purpose: To specify the cartridge on which the files are to be created.

Syntax: *SC , cart*

cart The CRN or –LU to be used for the scratch file.

Description:

Normally, FC builds each of its internal scratch files on the first cartridge with sufficient space for the initial size of the file. The SC command lets you specify the cartridge on which the files are to be created. (This command is useful as a way to avoid “Cartridge Full” errors in copy operations that require large amounts of scratch file space.)

You must enter the SC command before you perform a CO operation.

Title (TI)

Purpose: Sets the title to be included in the header file of tapes created by subsequent copy commands.

Syntax: *TI , title*

title The text, up to 72 ASCII characters, that specifies a title (extra characters are truncated).

Description:

Enter the title exactly as it is to appear in the header; lowercase characters are not converted to uppercase.

Transfer (TR)

Purpose: Allows you to transfer control to and return from a command file.

Syntax: *TR[, namr | –]*

namr Transfer control to the named command file.

– Transfer control to the terminal.

If neither *namr* or – are specified, transfer back to the next higher level command file when TR files are nested. When encountered in the first level command file, serves the same function as the FC Exit command.

Description:

FC can be run interactively from a terminal, from a command file, or as a combination of both.

Command files can be nested to four levels, with control passed from file level to file level using the TR,namr command. The current state of any of the configuration commands (Set List Device,

Title Header File, Name Comment File, Echo Commands, Group Commands) and the default values remain in effect through all levels of the nested command files.

Tape Handling

FC creates tapes in a unique format (shown at the end of this section), and thus cannot be written or read using the FMGR ST or DU commands. The program prompts you to mount the initial tape on the specified device, or to mount the next volume of a multi-volume set during the read or write operation, with one of the following messages:

```
put volume on LU n
```

```
put volume number n on LU n
```

FC checks the tape and device status and, if a problem is encountered, issues the appropriate error message:

```
tape LU n not ready
```

```
tape LU n media not initialized
```

```
tape LU n not write enabled
```

Each of the messages above is followed by one of these messages:

```
when ready to continue type GO, otherwise type BR
```

```
when ready to continue type GO, otherwise type BR or SK
```

If a “media not initialized” error occurs, initialize the CTD tape using the FORMC FO command (see the FORMC detailed description in this chapter for information on formatting).

When you enter BR in response to the “when ready” prompt, FC terminates cleanly.

The SK option is viable only if the operation is a multi-volume read. The SK response means skip the volume specified by *n* in the “put volume” message. In this case, FC repeats the “put volume” message for the next volume in the set, and reports “files lost” or “data lost” error messages for files on the SKipped volume. Refer to the section “FC Error Messages” later in this chapter for a description of the message text and corrective action.

FC automatically rewinds the tape at the beginning of a copy operation and writes over any existing data. For this reason, you cannot modify a tape created by FC; you must copy all files in a single operation. Use LH or DL to check the current contents before copying over a previously used tape. You can write a FMGR transfer file to run the FC LH command that issues a prompt before executing the copy. For example:

```
:RU,FC,LH,-8,K
:PA,,DO YOU WANT TO OVERWRITE THIS TAPE (type : ,YE or : ,NO)?
:IF,1G,EQ,YE,2
:IF,1G,NE,NO,-3
:TR
```

FC

:RUN, FC, CO, XX, -8
:TR

FC reads the first record of a tape volume and, if it does not conform to the FC tape format, either logs a tape format error or issues the message below (followed by the two records in ASCII):

```
tape not readable by FC, first two records are ...
```

FC locks the LU before a tape is accessed to prevent other programs from accessing the tape at the same time. If the tape is write-protected, the LU is unlocked at the end of the command. Otherwise, it remains locked until either FC encounters a tape-oriented command that specifies a different LU, a timeout occurs while FC is waiting for a command, or FC finishes executing.

To protect the tape, FC sets the device offline before the lock is released, unless you specified the K option in the last tape command or the command terminated due to an error or a break command. In those cases, the tape remains online. The CTD device is taken offline by unloading the cartridge (as if the UNLOAD button on the drive was pressed).

Since it can be inconvenient to have the device taken offline, you should use the write-protect mechanism or run FC interactively to retain the lock between commands. If you use the K option, which keeps the tape online in all circumstances, be aware that the tape is unprotected from access by other programs while it is unlocked and online.

If you interrupt a copy to tape by a BR (Break) command, the full tape generally is not usable. (In multi-volume sets, of course, volumes copied prior to the BR are intact.) However, if you specify BR during the verify phase, all data copied up to the point of the break is usable.

If you select so many files that one tape cannot hold the entire save, you must use multiple tapes. Large files are not copied across tape boundaries.

FC format on CTD:

Hdr	Comment	Dir.	EOF	data	data	EOF	...
-----	---------	------	-----	------	------	-----	-----

FC format on mag tape:

Hdr	EOF	Cmt fil	EOF	Dir.	fil	EOF	fil	EOF	...	EOF	EOF	...
-----	-----	---------	-----	------	-----	-----	-----	-----	-----	-----	-----	-----

Destination Disk Handling

FC performs a size check on destination cartridges on tape-to-disk copies before it copies any files. If it finds there is insufficient data or directory space for the files to be stored, FC issues the appropriate error message and exits. (Note that in open mode, other processes can create files on the cartridge; thus, it is possible for cartridge-full errors to occur even when the size check indicates sufficient space for the copy operation.)

For disk-to-disk copies, the size check is not performed and if a directory-full or cartridge-full error occurs, FC terminates and stores the status in the FMGR globals. (Refer to the section on “Error Handling in Transfer Files” later in this chapter for a definition of the FMGR global returns.)

Performance Considerations

The following steps can be taken to help increase the speed of the file copy operation.

1. Specify source cartridges explicitly if multiple cartridges are contained in the cartridge list. For example, if the cartridge list includes cartridges 10, 20 and 30 and you want to copy from cartridge 30, use a CO command as follows:

```
co, {a::30,b::30,c::30,d::30,e::30},#8
```

In this way, FC does not have to search cartridges 10 and 20 for the files to be copied. You may also use DE to specify the cartridge, as follows:

```
de, ::30
co, {a,b,c,d,e},-8
```

Note that in the command above, you may not put sequences of commands in the runstring.

2. Keep file extents collected. Copying speed is increased if the file main and its extents are adjacent on the source cartridge. To make the extents adjacent on all files on a cartridge, use a CO command, as follows:

```
co,100,,ldf
```

This copies the contents of cartridge 100 to itself (the destination cartridge defaults to the same cartridge as the source). You must specify the L option to ensure that the extents are adjacent. Select the D option so that the new files (with adjacent extents) can replace the old files. The F option logs the name of each file copied.

If a cartridge-full or directory-full error (FMGR-033 or FMGR-014) occurs during the above operation, pack the cartridge, using the FMGR PK command, and enter CO again. You may need to repeat this if the FMGR-033/-014 error recurs. If this happens on the same file, however, PK does not help, and you must use a larger cartridge for the operation.

If there are many small extents, reduce the number by making file sizes larger. To do this, copy each file with the ST command. Alternatively, you can eliminate all extents on a cartridge by using the CO command as shown and specify the E option. (Note that it is not always desirable to eliminate extents, because FMP is not able to re-use the space as easily when the files are purged, and larger file sizes tend to result in less efficient use of space.)

Extent collection (and elimination, if desired) can be done as a periodic cleanup function, much like a FMGR PK. If the extent collection/elimination process affects a lot of files (as shown by the F option log), issue a FMGR PK command after the CO command.

FC

3. Use Lock and Open modes efficiently. Determine which one works faster for an application, and use that mode whenever possible. Select the mode by specifying the L or O option in the CO command. Lock mode is more efficient for most copy operations that involve large numbers of files, because entire cartridges are locked in one operation, rather than by opening individual files. However, Lock mode may introduce other overhead, which could make it slower than Open mode. Open mode is usually more efficient when you are copying a small number of files. The faster mode for a given cartridge may depend on the following factors:
 - The number of files being copied to or from the cartridge.
 - The total number of files on the source disk cartridge.
 - Whether the destination disk cartridge was initially empty.
 - The number of files purged as a result of the D option.
4. Use cartridges efficiently. It is faster to copy to empty cartridges in Lock mode than to copy to cartridges containing other files.

Loading FC

Use the LINK command file #FCA to load the utility. The load process also requires copying the FC help file FCHLP to a cartridge available to all users.

Using Globals in Transfer Files

The information in the FMGR globals is helpful when you write transfer files to copy one source cartridge to a series of destination cartridges, moving to the next cartridge when the current cartridge is full. Such a file is particularly useful for backing up large, fixed disks to a series of flexible disks.

The transfer file makes use of the fact that the CO command file1 parameter specifies the first file to be copied, and the 10G global contains the name of the next file to be copied when a copy terminates due to a cartridge-full error. When you do a series of copies in which the ASCII contents of 10G are used as the file1 parameter, each copy can continue where the previous one stopped.

The following is an example of a transfer file used for this purpose. The source cartridge is specified in the 1G global, the negative flexible disk LU is specified in the 2G global, and the system master security code (msc) is specified in the 3G global. The msc is required for the ! option, which is specified to clear the destination disk without requiring operator intervention.

```
:PA,,Put in floppy and type TR
:MC,2G
:RU,FC,CO,1G,2G,!VF,,,3G
:DC,2G
:IF,5P,NE,-1,5
:PA,,Put in next floppy and type TR
```

```

:MC,2G
:RU,FC,CO,1G,2G,!VF,10G,,3G
:DC,2G
:IF,5P,EQ,-1,-5
:TR

```

Assuming a transfer file name FLOPUP, source CRN 100, destination LU 10, and msc XX, the following FMGR TR command causes CRN 100 to be backed up to a series of flexible disks on LU 10:

```
tr,flopup,100,-10,XX
```

If CRN 100 contains 25 files and 10 of these fit on the first flexible disk, 10 on the second, and five remain for the third, the process is:

```

:TR,FLOPUP,100,-10,XX
:PA,,Put in floppy and type TR
:TR
:MC,-10
:RU,FC,CO,100,-10,!VF,,,XX
FILE11                                (First ten files copied)
.
.
FILE11
FMGR-033                                (Cartridge-full error)
copy terminated
:DC,-10
:IF,-1,NE,-1,5
:PA,,Put in next floppy and type TR
:TR
:MC,-10
:RU,FC,CO,100,-10,!VF,FILE11,,XX
FILE11                                (Next file namr in 10G)
                                        (Next ten files copied)
.
.
FILE21
FMGR-033                                (Cartridge-full error)
copy terminated
:DC,-10
:IF,-1,EQ,-1,-5
:PA,,Put in next floppy and type TR
:TR
:MC,-10
:RU,FC,CO,100,-10,!VF,FILE21,,XX
FILE21                                (Next file namr in 10G)
                                        (Copy remaining files)
.
.
FILE25
:DC,-10
:IF,0,EQ,-1,-5                        (Test for cartridge-full error)
:TR
:                                        (Copy complete; exit)

```

FC

Error Handling in Transfer Files

Certain error conditions can be detected in FMGR transfer files by checking the FMGR globals after running FC, or, if FC is scheduled by a program other than FMGR, the scheduling program can retrieve the same information that appears in the FMGR globals by making a call to RMPAR. Since only one error can be returned in the globals, only the highest error number (the most severe error) is returned, even if multiple errors occur during FC execution. The parameters returned by RMPAR correspond to the FMGR globals as follows:

- PARAM(1) = 1P or characters 1-2 of 10G (next-source-file information)
- PARAM(2) = 2P or characters 3-4 of 10G
- PARAM(3) = 3P or characters 5-6 of 10G
- PARAM(4) = 4P (error category number)
- PARAM(5) = 5P (next-source-file information status)

4P can have the following values:

- 10000 – program aborted due to OF or system violation. Violations (such as Memory Protect) might result from a bad segment load.

The following errors terminate FC:

- 1000 – FC internal error
- 300 – transfer file stack overflow
- 200 – scratch file error (see SCRATCH command)
- 100 – segment loading error

The following errors terminate the current command:

- 90 – general error causing command termination
- 80 – destination disk cartridge or directory full
- 70 – break detected (system BR command)

The following errors allow the command to continue, but the purge phase is skipped on disk-to-tape copies:

- 50 – copying from current source volume aborted (tape-to-disk)
- 40 – unusual source file access error – file not selected
- 30 – general error on file being copied – file skipped
- 25 – tape format error (reading or verifying tape)

The following are warnings that do not affect FC operation:

- 20 – warning
- 10 – soft data error detected
- 0 – no error or warning

Global 5P can have the following values:

- 1 – next source file information not applicable
- 0 – no error
- 1 – next source file information valid (4P = 80)
- 2 – source is tape, number of next file greater than 32767

Globals 10G and 1P have possible values of:

If 5P = -1 (4P = 80) and source is a disk:
10G = name of next file to be copied from disk.

If 5P = -1 (4P = 80) and source is a tape:
1P = reference number of next file to be copied from tape.

If 5P = 1 and 4P not 10000, or 5P = -2, or 5P = 0:
10G = ASCII nulls
1P = 0

If 4P = 10000 (5P = 1):
10G and 1P are undefined.

In some circumstances, the meaning of the FMGR global information may be ambiguous. Since only the most severe error is posted to the FMGR globals, it is not possible to determine whether other errors occurred, or to determine the command that caused an error if multiple commands were entered. Furthermore, if a copy command has more than one source cartridge, the cartridge containing the “next source file” cannot be determined. If the GR copy command is used, additional ambiguity may result.

Because of this, the use of FMGR globals to determine the “next source file” is intended only for command files that contain a single CO command and involve a single source cartridge.

FC

FC Error Messages

This section contains a list of all the messages that FC generates. Brackets [] surround those parts of a message that are included only under certain conditions. Angle brackets < > indicate variables that are replaced with the appropriate value when the message is issued.

Errors Requiring Operator Action/Response

Do you really want to purge disk LU <#> CRN <crn> ?

A CO command specified the C option in the options parameter to clear the destination cartridge of all current data. FC requires a response to this prompt to be sure that you really want to destroy the data on the named LU or CRN. Use ! instead of C to clear the destination cartridge, if you want to suppress this warning prompt.

LU <#> is down correct problem and UP device, or use BR to break FC

Prior messages may help you understand the problem. Once the situation is corrected, bring the device up using the system UP command, and FC will continue. If you did not solve the problem before issuing UP, this message displays again. If you cannot correct the problem, issue BR to have FC terminate the current command.

tape LU <#> not ready when ready to continue type GO, otherwise type BR [or SK]

This error usually means that the magnetic tape device is not online or that the CTD cartridge is not loaded. Be sure that the correct LU was specified in the command. (Refer to the section on “Tape Handling,” earlier in this chapter, for an explanation of GO, BR, and SK.)

tape LU <#> media not initialized when ready to continue type GO, otherwise type BR

You tried to write to an uninitialized CTD cartridge. Use the FORMC FO (Format) command to initialize the cartridge. (Refer to the earlier section on “Tape Handling” for an explanation of GO and BR.)

tape LU <#> not write enabled when ready to continue type GO, otherwise type BR

For magnetic tape, you should install the write-protect ring to enable writing to the tape. For CTD, change the protect switch on the cartridge so that the arrow points away from the direction labeled “SAFE.” (Refer to the section on “Tape Handling” earlier in this chapter for an explanation of the GO and BR responses.)

put volume [number <#> on LU <tape LU> when ready to continue type GO, otherwise type BR [or SK]

The required tape should be mounted and online, or loaded. (Refer to the section on “Tape Handling” in this chapter for an explanation of the GO, BR, and SK responses.)

waiting to lock list device

FC tried to lock the list device, but discovered that either another program has the list device locked, or a resource number for the lock is unavailable. In either case, FC waits, retrying the lock every 2 seconds. Since operator intervention is not necessary, you can enter the system BR command to abort the command, rather than wait for the device or resource to become available.

Information Messages and Warnings

<source namr> [<dest namr>]

When a namr or a pair of namrs is logged with no other message, the indicated file was successfully copied from the source to the destination. When only a single namr is logged, it is the namr for the source file. When two namrs are logged, they are for the source and destination files, respectively. (The namrs may consist of the file name only, or may include the CRN.) Logging namrs of successfully copied files is done in Full mode only. (Refer to the “CO Command” section earlier in this chapter for a description of the Brief and Full modes.)

tape needed: <#> feet if 800 BPI, <#> feet if 1600 BPI

This message appears on a copy from disk to 1600 BPI magnetic tape if you specified the T option.

tape needed: <#> CTD blocks

This message appears on a copy from disk to CTD if you specified the T option.

scanning directories

FC is searching the cartridge directories to determine which files are being copied to tape.

copying files

FC is about to begin copying files to tape (assuming that the tape unit is ready).

cleaning up

FC finished copying files to tape, and is now closing source files or unlocking source cartridges. This process may take some time, and should be allowed to complete.

verifying volume

FC finished writing to the current volume and is now reading the volume in order to verify it.

break acknowledged

FC detected a system BR command, and is in the process of cleanly terminating the current command.

FC

writing tape at: <system time string>

Displays the date and time that are written in the tape header for this disk-to-tape copy.

beginning group

A GR command was entered, and subsequent CO commands are executed together once the EG (End Group) command is encountered.

group aborted

The AG (Abort Group) or AB (Abort FC) command was entered, causing all CO commands following the GR command to be ignored.

device is up, FC continuing

The device is now UP, and FC is attempting to continue the operation.

FC *xxxxx-xxxxx* REV.*xxxx* <*xxxxxx.xxxx*> Use ? for help.

This message appears before the first prompt is issued in interactive mode. The first line indicates the FC part number, the revision code, and the date/time code for the last revision.

(timeout)

FC detected a terminal timeout while waiting for input from the terminal. If a tape LU is locked, the lock is released, which may require taking the tape unit offline. Refer to the section on “Tape Handling” for details.

copy terminated

This message displays when a copy operation is terminated before completion for any reason, including an error or operator break.

title: <title>

Confirms the title entered in the TI command.

warning: title truncated

title: <title>

The title specified with TI was too long and was truncated. The truncated title is logged.

warning: no match for:[namr = <namr>][file1 = <?>][file2 = <?>]

No files were found that match both the source *namr* and the *file1* and *file2* parameters. Any combination of the three items may be listed, depending upon what was specified. Note that this does not mean that a match was not found for the *file1* or *file2* parameter itself, but that no files within the specified *file1/file2* range matched the *namr*.

warning: unable to eliminate extents in one or more files

Although the E option was specified, extents could not be eliminated from some files. See the description of the E option for details.

no files selected

The CO command was not executed, because no source files were selected.

disk write required retries:

LU <#> trk <#> sec <#> thru trk <#> sec <#>

The indicated disk write was unsuccessful on the first try, but succeeded after it was retried one or more times. Failure can be due to an error reported by the disk driver or a verify error, if the V option was specified.

disk read required retries:

LU <#> trk <#> sec <#> thru trk <#> sec <#>

Same as above, but for disk read.

disk directory write required retries:

LU <#> dir tr <#> thru dir tr <#>

Same as above, but for disk write in the directory area. The location of the read is isolated to a particular track or range of tracks. The track numbers indicated after “dir tr” are the directory track numbers, not the absolute track numbers on the disk LU. The first directory track (the one with the greatest absolute track number) is indicated as “dir tr 1”.

disk directory read required retries:

LU <#> dir tr <#> thru dir tr <#>

Same as above, but for disk read in the directory area.

tape format error <positive error code>

The data read from tape deviates from the expected format in some way, but copying from the tape can continue. The deviation is usually the result of a non-fatal tape read error. This message does not mean data is lost, but is provided primarily for diagnostic use. For tape format errors with a negative error code, see the section “Errors that Result in Rejection of Current Source Tape Volume”.

bad tape cartridge entry is:

<octal dump>

bad tape diskfile entry is:

<octal dump>

bad chunk header is:

<chunk size> <chunk type> <#entries> <1st file # or 0>

bad microdirectory entry is:

<file#> <file blk> <#blks> <buffer block>

The four messages above may accompany a tape format error message and are for diagnostic use only.

FC

I/O-*<error code>* on LU <#>[,D][,F]

This message has the same meaning as the operating system error message:

```
**I/O-error code @LU n[,D][,F]
```

See the *RTE-A User's Manual* for more information.

Disk Data I/O Errors

disk write failed: LU <#> trk <#> sec <#> thru trk <#> sec <#>

A disk write failed, even after retrying. The failure may be due to an error reported by the disk driver or a verify error, if the V option was specified. FC also reports a disk error (FMGR-001) for each file affected.

disk read failed: LU <#> trk <#> sec <#> thru trk <#> sec <#>

A disk read failed, even after retrying. The failure may be due to an error reported by the disk driver or a verify error, if the V option was specified.

FC also reports a disk error (FMGR-001) for each file affected. On disk-to-disk copies, the affected files are not copied. On disk-to-tape copies, the affected files are copied, but the affected parts of the files are flagged to indicate the disk error so that an appropriate error message can be given when the tape is read.

Non-Fatal Tape Read Errors

tape read error

The driver indicated an unrecoverable error on the tape read. FC does not retry in this case, because the driver would have retried if appropriate.

checksum error (chunk header)

The driver indicated a successful read, but FC's checksum information indicated a problem. FC tries again to recover; however, if retries fail, FC considers the information lost.

checksum error (chunk body)

The driver indicated a successful read, but FC's checksum information indicated a problem. FC tries again to recover; however, if retries fail, FC considers the information lost, unless the I option was specified. In that case, FC may try to use the information in spite of the checksum error, depending upon the context in which the error occurred. Appropriate warnings are displayed. See the "Data Lost" messages below for more information.

bad record length

The driver indicated a successful read (from magnetic tape), but the record length indicated by the transmission log returned from the driver disagreed with the length value specified in the header field of the record. FC tries to recover from this, but if retries fail, FC considers the information lost.

attempting to use data in spite of checksum error

See information above for “checksum error (chunk body).”

retrying tape read

A “checksum error ...” or “read record length” error occurred, and FC is attempting to recover from the error by retrying the tape read.

Errors Affecting a Single File**can't open <name>::FMGR <nnn>**

The indicated source file could not be opened to copy. Typically, this is because the file is open exclusively to another program (FMGR-008), or because the cartridge is locked to another program (FMGR-013).

file <name>::purged or replaced during copy or directory entry corrupt

This file was purged and possibly replaced by a different file when FC was scanning the directory, or a directory entry for the indicated file is corrupt. The file is not copied.

**<source namr> [<dest namr>]
FMGR <nnn>**

A namr or pair of namrs followed by a FMGR error code means the file was not copied because of the FMGR error. However, for a FMGR-001 error on a disk-to-tape copy, the file is copied to tape but the affected parts of the file are flagged to indicate the disk error, so that an appropriate error message can be given when the tape is read.

When a single namr is logged, it is the namr for the source file. When a pair of namrs are logged, they are for the source and destination files, respectively. Namrs may consist of the file name only, or may include the CRN.

**<source namr> [<dest namr>]
FMGR <nnn>
warning: above file copied but not purged from source**

Similar to the above error, except that the FMGR error only prevented the source file from being purged as requested by the P option; the file was still copied.

FC

**<source namr> [<dest namr>]
file too large for this operating system**

Similar to the above error, except that the “file too large ...” error occurred, rather than a FMGR error. The file cannot be copied because it is larger than 16383 blocks. This may happen when you copy files from a tape created under another operating system that supports files of that size.

**<source namr> [<dest namr>]
data lost**

Similar to the above error, except that the “data lost” error occurred rather than an FMGR error. Data lost errors can only occur on tape-to-disk copies, and indicate that file data was lost due to a problem reading the source tape. A read problem may be identified in earlier messages (see “Non-fatal Tape Read Errors”), or a record may have been dropped, in which case no problem is detected other than the data lost error itself.

For data lost errors, other messages may intervene between the line with the namrs and the data lost message. The data lost error applies to the most recently logged file namrs, even if the namrs do not immediately precede the data lost message.

If the message “checksum error (chunk body)” appears anywhere before a “data lost” message, you may recover some of the lost data by repeating the tape-to-disk copy with the I option specified. When I is used, destination files are created even though some or all of the data was lost, and “data lost” messages are replaced by some of the more specific messages below, which indicate the exact part of the file affected and whether the data loss is certain.

**<source namr> [<dest namr>]
data lost due to disk error when tape was made**

Same as above, except the data was lost because of a disk error that occurred when the tape was written, not because of a problem reading the tape. An FMGR-001 (disk error) occurred when this file was copied to tape.

**<source namr> [<dest namr>]
data lost: entire file**

Same as “data lost,” except this message occurs when the I option was used to try to recover part of the file, but the attempt did not succeed. Because I was specified, the destination file is created even though all of its data was lost. The resulting destination file is of no value, except for identifying information that may be obtained from the directory entry.

<source namr> [<dest namr>]

data lost:

[xtnt <x1>] blk <b1> thru [xtnt <xn>] blk <bn>

Same as above, except that only part of the file was lost. The message indicates the range of blocks that were lost. The first and last block in the range are specified by giving the extent number (omitted if the block is in the main extent) and the relative block number within the main or extent. Blocks are numbered from zero within the main or extent.

More than one of these messages may be given for a single file. The file namrs are not repeated for subsequent data lost messages on the same file, even though other messages may intervene, thereby breaking up the sequence of data lost messages.

Because the I option was specified, the destination file is created even though some or all of its data was lost.

<source namr> [<dest namr>]

possible data loss: entire file

Same as “data lost: entire file” except that it is uncertain whether any data was lost. The file should be examined to determine the loss, if any.

<source namr> [<dest namr>]

possible data loss:

[xtnt <x1>] blk <b1> thru [xtnt <xn>] blk <bn>

Same as above, except that possibly only part of the file was lost. The message indicates the range of blocks that may be lost. The first and last block in the range are specified by giving the extent number (omitted if the block is in the main extent) and the relative block number within the main or extent. Blocks are numbered from zero within the main or extent.

More than one of these messages may be given for a single file. The file namrs are not repeated for subsequent data lost messages on the same file, even though other messages may intervene, thereby breaking up the sequence of data lost messages.

<source namr> [<dest namr>]

data lost due to disk error when tape was made: entire file

Same as “data lost due to disk error when tape was made,” except this message occurs when the I option was used to try to recover part of the file, and the attempt did not succeed.

Because the I option was specified, the destination file is created even though all of its data was lost. The resulting destination file is of no value, except for any identifying information that may be obtained from the directory entry.

FC

**<source namr> [<dest namr>]
data lost due to disk error when tape was made:
[xtnt <x1>] blk <b1> thru [xtnt <xn>] blk <bn>**

Same as above, except that only part of the file was lost. The message indicates the range of blocks that were lost. The first and last block in the range are specified by giving the extent number (omitted if the block is in the main extent) and the relative block number within the main or extent. Blocks are numbered starting at zero within the main or extent.

More than one of these messages may be given for a single file. The file namrs are not repeated for subsequent data lost messages on the same file, even though other messages may intervene, breaking up the sequence of data lost messages. This message may appear repeatedly for different parts of the same file, even if the different parts are adjacent and even if the net meaning of all the messages is that the whole file was lost.

Because the I option was specified, the destination file is created even though some or all of its data may have been lost.

source file <name>::selected by commands with conflicting parameters

On a tape-to-disk copy specified by a GR command with more than one CO command, the indicated source file was selected by more than one CO command, but those particular commands did not meet the following restrictions:

- The destination parameters must be the same in all the commands.
- The D option must be consistently selected or not selected in all the commands.
- The E option must be consistently selected or not selected in all the commands.
- The msc parameter must be consistent through all the commands.
- You must not specify the L option in some commands and the O option in others.

Loss of Unidentified Files on Copy from Tape

files lost, reference numbers <n1> thru <n2> names not available

Indicates that distributed directory information was lost for files with reference numbers n1 through n2. Because this information was lost, the names and other information about the files is not known. You can list the main tape directory with the DL command to determine what files these numbers correspond to.

Disk-to-Tape Copy Verify Errors

files lost, reference numbers <n1> thru <n2>

Indicates that distributed directory information was lost for files with reference numbers n1 through n2. If necessary, you can list the main tape directory with the DL command to determine what files these numbers correspond to.

data lost for file, reference number = <#>

Indicates that data was lost for file reference number n. If necessary, you may list the main tape directory with the DL command to determine what file this number corresponds to.

error: header file verify error

Usually indicates that no data was actually written to the tape, probably due to a defective write head on the tape drive.

Errors that Cause Rejection of Current Source Tape Volume

tape not readable by FC, first two records are:

<record 1>

<record 2>

The tape mounted is apparently not an FC tape. The first two records are logged to help determine what the tape is.

volume does not match others

On a multi-volume tape-to-disk copy, the tape volume just mounted does not come from the same set as the previous volumes.

wrong volume

The tape volume just mounted does not have the correct volume number. Volumes in a multi-volume tape set must be restored in the correct sequence. When FC prompts for a new volume, it indicates which volume number is to be mounted.

fatal tape read error

The tape driver reported an error when the header or comment file was read. FC makes no attempt to recover from this; however, you can retry the operation manually. On multi-volume tape-to-disk copies, you need not start over with the first volume; just enter GO at the prompt to mount the volume. Refer to the earlier section in this chapter, "Tape Handling," for an explanation of the options available at this point.

header file checksum error

The information in the header file is questionable, due to a checksum error. It would be unwise for FC to continue reading the volume in this situation. You may retry the operation manually.

FC

tape format error <*negative error code*>

This error indicates that the data read from tape deviated from the expected format so severely that FC cannot continue reading the volume. For CTD, this usually means that the tape is blank or is not a CTD tape; otherwise, it is probably the result of a non-fatal tape-read error. The negative error code is for diagnostic use only. You may retry the operation manually.

For tape format errors with positive error codes, see the section on “Information and Warning Messages.”

Command Syntax, Parameter Errors (Command is Skipped)

error: no such command (use ? for help)

error: bad or missing parameter

error: braces used where not permitted

These error messages are self-explanatory.

error: bad namr

An incorrectly formed namr appeared in the command.

error: mismatched braces

Self-explanatory.

error: namr list not allowed in this context

A list of namrs delimited by braces was used in a parameter where a list of namrs is not permitted.

error: unrecognized option character

An unrecognized character appeared in the “options” parameter.

error: incompatible options

Contradictory options, such as B and F, or L and O, were specified in the same command.

error: option not defined for this command

An option character that is not meaningful for this command was specified in the “options” parameter.

error: option not applicable

An option character that is not applicable for this type of copy (disk-to-disk, disk-to-tape, or tape-to-disk) was specified in a CO command.

error: bad file1 or file2 param

The value for the file1 or file2 parameter was incorrectly specified. These parameters specify file names when copying from disk, and file numbers when copying from tape.

error: bad msc parameter

Self-explanatory.

error: no tape-to-tape copy

Both the source and destination parameter specified a tape device, implying a tape-to-tape copy, which is not supported. Make sure you have not used the DE command incorrectly.

error: source or destination incompatible with group

The source and destination parameters must be either a disk or a tape consistently throughout the grouped commands, and you may only specify one tape LU in a given group.

error: options inconsistent with group

Within a group of commands, you may not specify certain options--V, I, S, F, B, T, and K--in one CO command, but omit them from another. DE may be useful to ensure these options are used consistently in all the CO commands within a group.

error: L and O options must be consistent for a given cartridge

The L and O options must be consistent for a given disk cartridge within a group. Therefore, both L and O may not be specified for the same cartridge, even in different CO commands within a group. You may specify L or O in one command, and no option in another. (In this case, the specified option affects both commands.)

error: renaming multiple files to same name, or cartridge not found

The source name was omitted, or has wildcards, or the source is a list of namrs, but the destination specifies a file name. A destination file name is permitted only if it is clear that the source parameter can select only one file.

It is possible that the destination parameter was intended to specify an ASCII CRN in the abbreviated form (CRN rather than ::CRN), but the ASCII CRN was interpreted as a file name because no cartridge with that CRN was found on the cartridge list.

error: bad destination name

An illegal filename was specified in the destination parameter.

error: filename BOOTEX reserved for system

The destination parameter specified BOOTEX as the destination file name. This is not permitted, because the name BOOTEX is reserved for the system boot extension file.

error: on disk-to-tape copy, dest cartridge may not be specified as -LU

If you specify a cartridge in the destination parameter on a disk-to-tape copy, it must be a positive CRN, not a negative LU.

FC

error: P option requires explicit source cartridge

You may not specify the P option without also specifying the source disk cartridge explicitly in the source parameter.

error: L and O options require explicit source cartridge

You may not specify the L or O option without also specifying the source disk cartridge explicitly in the source parameter.

error: P option not allowed if source and dest are same cartridge

You may not specify the P option on a disk-to-disk copy in which the source and destination are on the same cartridge.

error: C and ! options require explicit destination cartridge

When you specify the C or ! option, you must also specify the destination disk cartridge explicitly in the destination parameter.

error: dest secu code not allowed w/o source secu code or good msc

A security code was specified in the destination parameter, but no security code was specified in the source parameter and the msc parameter was not specified correctly.

error: bad LU or device not supported

An LU specified in the command is not the LU of a disk or tape device supported by FC.

error: bad tape LU

The tape LU specified in the command is not the LU of a tape device supported by FC.

Command Out-of-Sequence Errors (Command is Skipped)

error: already in group

A GR command was entered before the previous GR ended.

error: not in group

An EG or AG command was entered, but there was no previous GR to end or abort.

Other Errors that Terminate Current Command

error on help file FCHLP: FMGR <nnn>

The indicated error occurred during an attempt to access FC's help file, FCHLP, needed for the ? command.

cartridge lock error on disk LU <#>: FMGR <nnn>

FC was unable to lock the indicated disk cartridge due to the FMGR error. If the error is FMGR-103 (Corrupt Directory Detected), resolve the problem before trying to access any files on the cartridge, otherwise unnecessary loss of data may result.

cartridge unlock error on disk LU <#>: FMGR <nnn>

Same as above, but for cartridge unlock.

cartridge clearing error on disk LU <#>: FMGR <nnn>

Same as above, but for clearing the directory on the disk cartridge.

error in cartridge status on disk LU <#>: FMGR <nnn>

Same as above, but more general.

disk directory write failed: LU <#> dir tr <#> [thru dir tr <#> error writing directory on disk LU <#>: FMGR <nnn>

A directory write to the specified track or tracks failed due to a disk error (FMGR-001) or verify error (FMGR-049), even after retries were attempted. The location of the write is isolated to a particular track or range of tracks. The track numbers indicated after "dir tr" are the directory track numbers, not the absolute track numbers on the disk LU. The first directory track (the one with the greatest absolute track number) is indicated as "dir tr 1".

disk directory read failed: LU <#> dir tr <#> [thru dir tr <#>] error reading directory on disk LU <#>: FMGR <nnn>

Same as above, but for directory read instead of write.

comment file error: FMGR <nnn>

The indicated FMGR error occurred when the comment file (determined by the most recent CF command) was opened or read.

FC

error: number of cartridges to be copied to tape exceeds limit

Cartridges were renamed on a copy to tape, resulting in the tape's cartridge list containing more than 64 distinct cartridges.

error: volume not big enough for header/comment/directory files

The destination tape is too small to be used. There is not enough room for the minimum amount of information that must be written before going on to another volume.

error: fatal i/o error on LU <#> status= <octal A-Register status>

FC attempted I/O to the specified LU and the status returned in the A-Register indicates a problem that should never occur. The status value in the message is the value returned in the A-Register from the EXEC call.

error: comment file checksum error

A checksum error was detected when reading the comment file from the tape.

error: not enough memory

FC does not have enough free memory to let the command proceed. If FC was sized to fewer than 32 pages, increasing the size may solve the problem. If you used the GR command, try using fewer CO commands in the group to prevent this problem.

error: tape lock failed

FC cannot lock the tape LU required by the command. Either the tape LU is locked to another program, or no resource number is available for the lock.

error: tape DESCRIBE error

The CTD driver reported an error when FC issued a "DESCRIBE" control request.

error: error on tape write

An unrecoverable error was reported by the driver on a tape write request. FC does not retry in this case, because the driver would have retried if appropriate. FC cannot continue copying to tape after this, as the resulting tape would be corrupt.

CRN <CRN> not found

A tape-to-disk CO command with an unspecified destination cartridge was issued, and the required CRN, determined by reading the tape directory, was not found.

CRN <CRN> LU <#> not large enough: data blocks/dir entries needed: <bl>/<ent> available: <bl>/<ent>

The specified cartridge does not have enough file or directory space to copy the files from the tape.

Other Errors

list file error: FMGR <nnn>

An error occurred when FC attempted to access the list file specified with the most recent LL command. This causes the list device to revert to the terminal; it does not cause the command to terminate.

command file error: FMGR <nnn>

An error occurred when FC tried to read a record from the command file (determined by the most recent TR command). This causes an automatic TR to the terminal (FC enters interactive mode).

Errors that Terminate FC

fatal scratch file error: FMGR <nnn>

An error occurred when FC tried to create or access one of its scratch files. Each scratch file is created on the first cartridge that has room for the initial size of the file, unless you specified a cartridge in the SC command.

If the error code is FMGR-033 (disk cartridge full) or FMGR-014 (directory full), make more room on the cartridge or use SC to have FC create the scratch files on a cartridge with more room.

If the error code is FMGR-012 (EOF or SOF error), you probably included too many CO commands in a group. The maximum number of grouped copy commands is approximately 100. (A command that has more than one namr in the source parameter must be counted as a separate command for each occurrence of a namr.)

An error (FMGR-011) can occur when you copy more than 5000 files (approximately) from disk to tape in Lock mode. The error only occurs, however, if FC creates a scratch file on a cartridge that is being locked for the copy. This error occurs during the “scanning directories” phase of the disk-to-tape copy, when no files are as yet copied to tape.

To avoid this problem when you need to copy more than 5000 files to tape, use the FC command, SC, to specify a scratch cartridge other than one from which files are being copied. If that is not possible, copy files to tape using the Open mode, rather than Lock, by specifying the O option in the CO command.

For any other error, if SC is used, make sure it is executed before any CO commands.

can't load segment <segment name>

The indicated segment could not be loaded.

FC

TR stack overflow

The command file stack used for the TR command overflowed. This error occurs if more than four nested TR commands are given (excluding any TR command in the runstring).

internal error at *<address>* last segment loaded: *<segment name>*

FC detected an inconsistency that cannot be interpreted. This error should never occur under normal conditions. The address and segment name are for diagnostic use only.

internal error: pascal *<error type#>* *<error#>* *<line#>*

FC is being terminated due to a Pascal run-time error. This error should never occur under normal conditions. The numeric values are for diagnostic use only.

exec error *<system error mnemonic>* exec params: *<params in octal>*

The operating system or I/O driver rejected an EXEC call made by FC. This error should never occur under normal conditions. The numeric values are for diagnostic use only.

FORMA

FORMA provides initialize and spare functions for Micro/1000 systems with internal disks connected via the HP 12022 disk controller. A series of prompts helps you select a function and further configure the program to your specific application.

FORMA is designed to work together with ERTSH, an error rate test program for mini-Winchester hard disks. ERTSH is discussed in detail in the section following FORMA. Note that FORMA and ERTSH require at least 256 KBytes to run.

A major cause of disk problems is minute flaws in the magnetic disk surface that cause some data bits to be unreadable or to be read incorrectly. Because these defects are on the physical disk surface, they may not be centered on a track and may cross logical boundaries. It is common to find a sector with a correctable error next to one that cannot be corrected. You can use FORMA to spare the correctable error before it grows into an error that cannot be fixed.

To spare bad tracks, the disk must be in FORMA format. If your disk is currently in FORMF format, you should re-initialize it, using FORMA, at your earliest convenience (for example, after your next backup). If you wait until major errors appear, the defects could be severe enough to prevent a successful backup.

Using ERTSH should become a regular part of your periodic maintenance procedure for the disk. The ERTSH Read Only (RO) command indicates the general state of the disk without overwriting data. Other combinations of commands discussed in this chapter let you test for and automatically spare defective tracks on disks that are in FORMA format.

FORMA Definitions

The following definitions apply to formatting and sparing Micro/1000 internal mini-Winchester disks.

Index pulse	A physical marker on the spindle of the mechanism that generates a logical signal once every revolution. It is used to designate the beginning of a track. See Figure 2-3.
Primary spare	A spare that covers a defect that was originally designated by the manufacturer. All other spares are called secondary spares, including any spares found after shipment from HP. Both types of spares are created using the FORMA SS (Spare Sector) command.
Dormant spare	A reserved sector (sector 32) that is used for sparing if a defect appears on that track.
Track format	The logical layout of a section of the disk at one head and cylinder that provides consistent and efficient data transfer.
ID field	A non-data section that contains addressing information for that sector. See Figure 2-3. It is read by the disk controller before each data field is read in allowing sectors to be in any order on a track. The process of writing these fields on every sector is part of the IM (Initialize Media) command.

FORMA

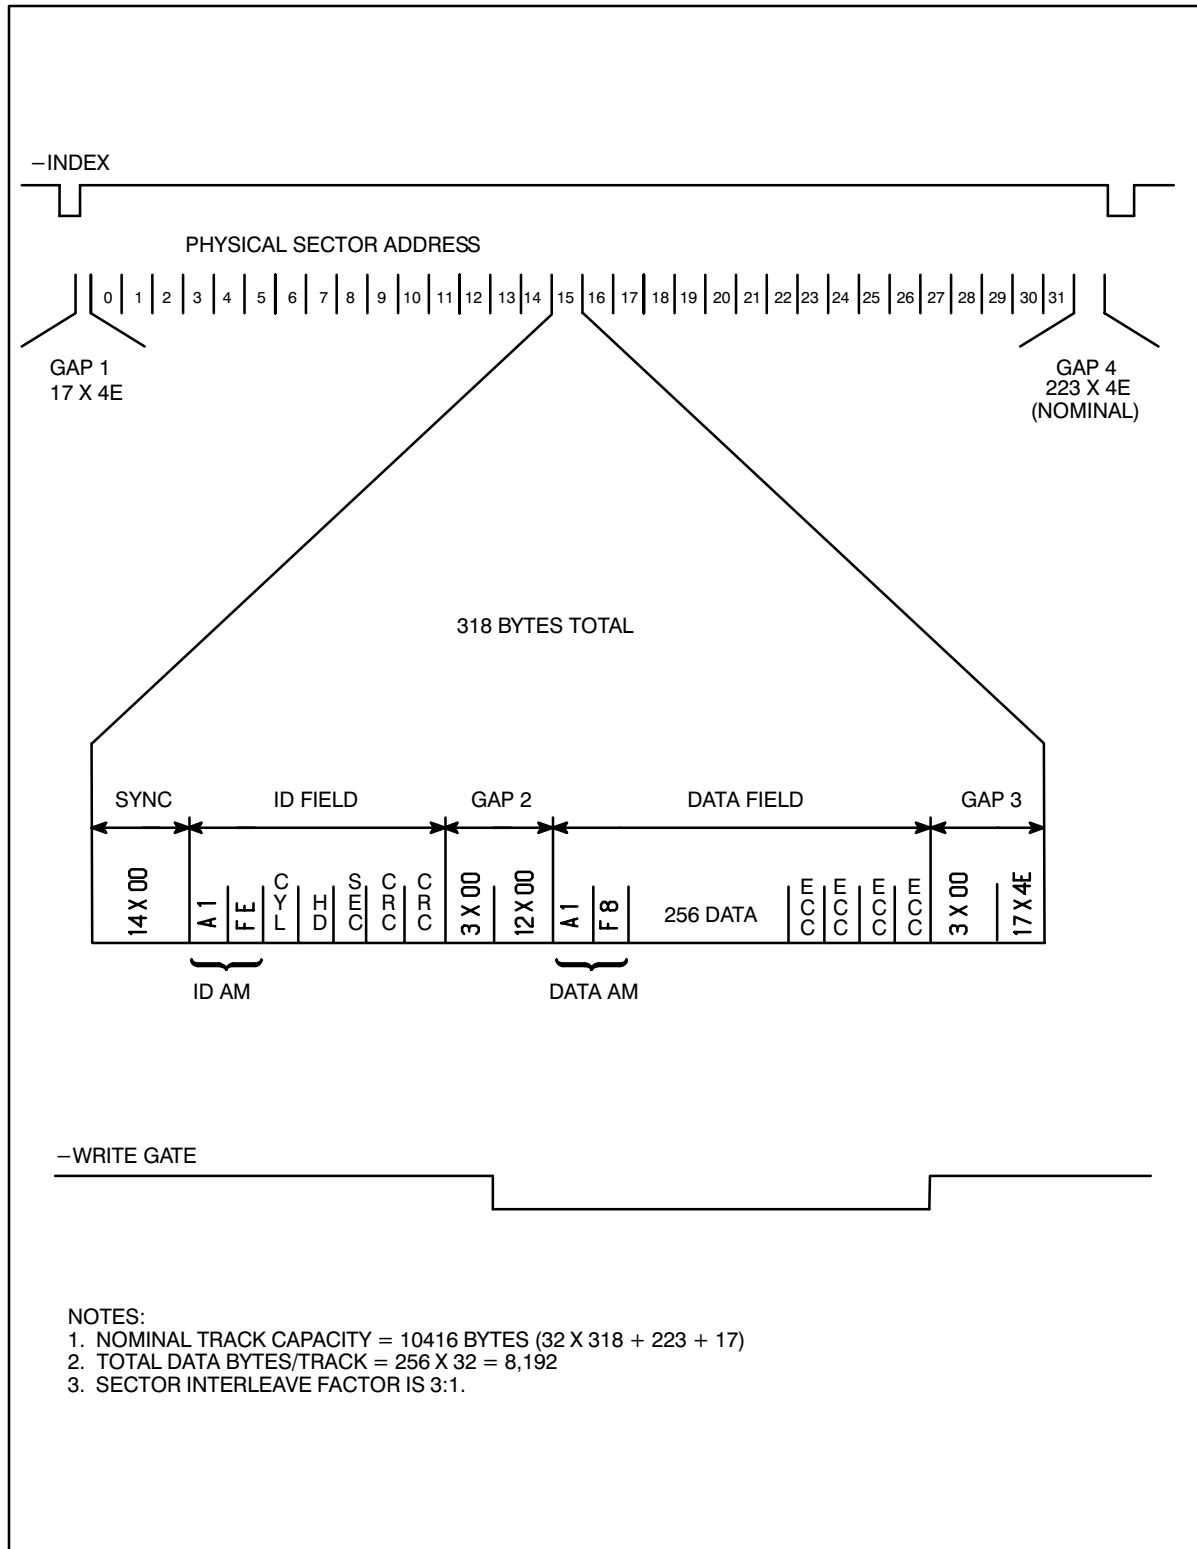


Figure 2-3. Typical Track Format

Calling FORMA

To call FORMA, enter the following runstring:

```
CI> [RU, ]FORMA ,command ,disk lu ,option
```

The command and option parameters are discussed in the following sections.

Alternatively, you can run FORMA interactively by entering the following:

```
CI> [RU, ]FORMA
```

The program prompts you for the commands, LU, and parameters.

FORMA Commands

Formatting and sparing functions are carried out by two FORMA commands, IM (Initialize Media) and SS (Spare Sector). The IM command is used during installation or for recovery from major damage to media. The SS command spares defective sectors by recovering data and assigning a spare track that the disk controller can access.

Help (??)

Purpose: Enter this command at the TASK? prompt to see a list of available responses.

Syntax: ??

Initialize Media (IM)

Purpose: Initializes a new disk when it is installed, re-initializes a disk previously initialized with FORMF, and recovers from major disk media damage.

Syntax: IM [*lu*] [*spare_option*]

lu The disk logical unit.

spare_option The option to save primary (P) spares, secondary and primary (S) spares, or no (N) spares.

Description:

IM provides three options, P, S, and N, which you can choose to preserve primary, secondary and primary, or no spares. Always retain primary spares, unless the media is severely damaged.

FORMA

Caution FORMA initializes an entire disk upon recognition of any valid LU number. Before initializing a disk that is currently in use, be sure to back up the disk.

New disks shipped from HP are already initialized and spared in FORMA format; you need not initialize them again. In this case, proceed to the discussion of ERTSH later in this chapter.

FORMA and FORMF formats are not compatible. Using FORMA to re-initialize disks that were previously initialized with FORMF takes a long time (up to 16 hours). If the floppy disks were initialized with FORMF, do not retain any spares before you re-initialize with FORMA.

Do not use IM in place of a logical debugging process; you should only use it at installation or in case of total media loss.

Spare Sector (SS)

Purpose: Spares a sector to an unused portion of the disk reserved for that purpose. Retention of data from the sectors being spared can be specified.

Syntax: `SS , <disk lu> , <r> , <s | b> , <ccc> , <h> , <ss | bbbb> , [p] , [o | u]`

<i>disk lu</i>	The number of the LU, from 1 to 63.
<i>r</i>	The option to retain data (y or n).
<i>s</i>	Logical (sector) addressing.
<i>b</i>	Physical (byte) addressing.
<i>ccc</i>	The cylinder number of the defect.
<i>h</i>	The head number.
<i>ss</i>	The logical sector since index (see index description below).
<i>bbbb</i>	The physical byte count since index (see index description below).
<i>p</i>	Save the primary spare.
<i>o</i>	Override the old spare.
<i>u</i>	The “unspare” option.

Description:

You must enter the first six parameters or the program aborts. (The options are further described in the next section.)

Note that when you use SS, there is some chance of data loss on the spared sector, depending upon the number of defective bits in that sector.

You must enter only the byte/index address of the defective sector when performing sparing. Logical sectors refer to the address written on the ID field of the sector during formatting. Logical sectors with adjacent addresses are not physically adjacent, due to a 3:1 interleave. Table 2-16 provides information that you need to answer the system prompts.

Table 2-16. Responses to System Prompts

	Cylinder	Head	Sector	Byte/Index
10Mb	0 – 305	0 – 3	0 – 30	16 – 10048
15Mb	0 – 305	0 – 5	0 – 30	16 – 10048
20Mb	0 – 611	0 – 3	0 – 30	16 – 10048

Byte/index refers to the number of bytes that passed under the head since the last index pulse (which signifies the physical start of a track). This form of input is useful when you are sparing sectors that the manufacturer has indicated are defective. Use the label on the drive mechanism, as discussed in the section in this chapter on “Installing a New Disk.” The label uses the byte/index format. One spare sector per track is available.

Remember that some defects will fall into gaps between the data fields. FORMA detects these and leaves them unspared, returning this message:

```
FORMA: REQUESTED SPARE FALLS WITHIN A NON DATA AREA, NO SPARE REQUIRED
```

SS Command Options

When you enter the P (Primary Spare) option as the seventh parameter, a primary spare is allocated on the track specified. All spares made from the information derived from the label on the drive should be primary spares. (See the section “Installing a New Disk” in this chapter.)

You can use the O (Override Spare) and U (Unspare) options to correct errors made when sparing. Override allows sparing when a sector is already spared on a track. Unspare re-initializes a track to the standard dormant spare state.

Converting from FORMF to FORMA

Although you can run the error rate test (ERTSH) on disks that are in FORMF format, you must use FORMA to initialize a disk before you can have defective tracks spared. The FORMA initialization procedure is similar to the installation of a new disk, described in the following section, except that references to the manufacturer’s label do not apply.

Installing a New Disk

Follow these steps to install a new internal hard disk:

1. Find the label attached to the top of the disk drive that lists the areas the manufacturer has found to be defective. Record the CYL, HD, and BYTE/INDEX information. This data is used later in the FORMA runstring.
2. If you already installed the disk, your customer engineer (CE) can read the label by removing the three securing screws from the bottom of the peripheral tray and sliding the disk drive out far enough to read the label. If you are working without the aid of an HP CE, do not attempt to remove the disk drive. Use ERTSH to locate the defects. ERTSH is discussed in detail later in this chapter.

FORMA

FORMA Error Handling

All user errors cause FORMA to abort before execution. The program reports disk errors during execution of a disk operation. The current status of any disk error is always printed to the terminal.

FORMA Error Messages

Error messages fall into three categories: user messages, user errors, and disk errors.

User Messages

FORMA: INITIALIZE COMPLETE

This message displays when IM execution completes.

FORMA: OVERRIDE SPARE COMPLETE

This message displays when you specify the Override option.

FORMA: RESTORING DATA

After sparing, the whole track is re-written to the disk.

FORMA: SAVING SECTOR XX

This message indicates each sector being spared.

FORMA: SPARING COMPLETE

The standard spare operation is complete.

FORMA: UNSPARE COMPLETE

This message displays when you specify the Unspare option.

User Errors

The following messages print at the terminal when FORMA encounters a non-standard condition.

FORMA: ILLEGAL COMMAND

You entered an unsupported mnemonic.

FORMA: ILLEGAL DISK LU

Either the logical unit specified is not generated as an integrated hard disk, or the wrong LU was entered. FORMA does not accept any microfloppy disk LUs or non-integrated disk LUs.

FORMA: ILLEGAL PARAMETER

A parameter was left out or the wrong values were inserted in the fields.

FORMA: ILLEGAL RESPONSE

An interactive mode prompt was answered with an illegal type, either alpha when numeric was requested, or the reverse. The prompt is repeated.

FORMA: ILLEGAL USE OF OVERRIDE FIELD

You specified something other than U or O in the last field.

FORMA: INSUFFICIENT CAPABILITY

The user is not a superuser.

FORMA: MISSING PARAMETER

You entered too few parameters in the runstring.

FORMA: NOT AN OPTION

You entered characters other than SS or IM as a command.

Disk Errors

FORMA: CONTROLLER FAULT

The controller has passed a status of 11400B, indicating a possible hardware problem.

FORMA: DRIVE NOT READY

The disk drive does not respond to the first IO request. The problem may be related to the hardware configuration, the generation IO configuration, or an incorrect configuration of the shunt block on the mechanism.

FORMA: FATAL INTERNAL ERROR

The status from the HP 12022 controller does not match any defined status. Check the ribbon cables, processor/memory controller, or other hardware.

FORMA

Disk Error Decoding

Because the disk, disk controller, and interface card are all involved in disk reads and writes, running FORMA could result in the disk interface driver (ID*27) detecting an error. When illegal status is detected, this error code is displayed (when in runstring mode) as follows:

FORMA: DVT16 = xxxxxB DVT17 = xxxxxB DVT18 = xxxxxB DVT19 = xxxxxB

This is the status returned by the HP 12022A controller when an error condition occurs. Use the following chart to decode the information.

<u>DVT16</u>	<u>DVT18</u>	<u>DVT19</u>	
1	0	0	Request Rejected by Driver (Set by Reject)
2	0	100002B	Drive Not Ready or Seek Not Done
3	0	0	Timeout on DMA
5	0	0	Parity Error on Transmission
10B	0	0	Corrected Data Error
12B	1400B	0	Configuration Switches Indicate No Drive is Connected, Set by No Drive
12B	11400B	100020B	Busy, Write Fault, TR00 Error, No Handshake Completion
100077B	4400B	0	Bad Block or ECC/CRC Error
100077B	4400B	100004B	ID Not Found or DAM Not Found

DVT17 always contains 0.

For more information on the operation of Driver ID*27, see the *RTE-A Driver Reference Manual*, part number 92077-90011.

FORMC

FORMC is an online maintenance program for CS/80 disk and tape drives. You can use it to format an entire disk volume or tape, verify the integrity of data on the device, and correct any errors detected. FORMC deals only with logical addresses; for offline use or if physical addressing of a device is desired, use the offline version of FORMC (!FORMC) or the CS/80 disk exerciser program, EXER. Refer to the *CS/80 Disk Exerciser Manual*, part number 5955-3462, or Appendix C, for details on the EXER utility.

Calling FORMC

You may run FORMC interactively or in runstring mode, as described below. FORMC can also be called programmatically by issuing the appropriate program-scheduling EXEC call and entering the parameters in the runstring.

To call FORMC, enter the following runstring:

```
CI> [RU, ]FORMC[ ,list LU] ,command ,lu [ ,params ]
```

where:

<i>list LU</i>	The LU of the list device to which FORMC directs messages generated during the specified operation.
<i>command</i>	One of the FORMC commands described below.
<i>lu</i>	The LU of the CS/80 disk or tape to be accessed by FORMC for the specified operation.
<i>params</i>	The command descriptions are provided below.

If you omit a required parameter or enter one incorrectly, FORMC exits with the messages:

```
/FORMC: ILLEGAL RUNSTRING PARAMETER
/FORMC: ABORTED
```

If you do not enter any runstring parameters, FORMC runs interactively and prompts you for the commands, as follows:

```
/FORMC: CS80 DISK FORMAT UTILITY
/FORMC: '?' WILL LIST THE LEGAL COMMANDS
/FORMC: TASK?
```

When you enter a command, FORMC prompts for parameters.

FORMC

Command Execution

Before it executes a command, FORMC issues the message:

```
/FORMC:  COMMAND EXECUTING - <command> <parameters>
```

FORMC terminates after successful command execution or, if an error occurs, after it issues the appropriate error message.

Device Driver Status

During command execution, FORMC tests the driver status and if -1 is present in QSTAT (indicating driver timeout), it issues the following messages and exits:

```
/FORMC:  DEVICE DRIVER TIMED OUT. CONTACT SYSTEM MANAGER.  
/FORMC:  ABORTED
```

Driver timeout can result if the disk drive is disconnected or if a hardware failure occurs. Refer to the “FORMC Error Messages” section later in this chapter for a description of the CS/80 device driver status information message.

Break Detection

FORMC checks the break flag only during verification after each locate and read. If the flag is set, FORMC exits with the following messages:

```
/FORMC:  BREAK DETECTED  
/FORMC:  ABORTED
```

The break flag is not checked during formatting and sparing operations. To protect data integrity, these operations should complete without intervention after they are initiated.

FORMC Commands

FORMC commands are summarized in Table 2-17 below and described in the sections that follow.

Table 2-17. FORMC Commands Summary

Commands	Description
Information Command	
HElp ?	Displays a list of FORMC commands
Configuration Commands	
FO rmat <i>disk tape</i>	Formats the specified disk or tape
SP are <i>disk track</i>	Spares the specified disk track
VE rify <i>disk tape</i>	Verifies the integrity of the specified disk or tape
Exit Commands	
AB ort EN d EX it	Any of these commands causes FORMC to terminate

Abort, End, and Exit (AB) (EN) (EX) (/E)

Purpose: To abort or terminate FORMC.

Syntax: AB, EN, EX, or /E

Description:

The utility stops execution when you enter any of these commands in response to any FORMC prompt.

Caution DO NOT abort the FORMC program during a spare or format operation, as you may jeopardize the integrity of the disk data.

FORMC

Format (FO)

Purpose: Formats a disk volume or cartridge tape.

Syntax: `FO, media_LU [,interleave]`

media_LU The LU of the CS/80 disk or tape to be formatted. FORMC aborts with an ILLEGAL RUNSTRING PARAMETER message if you try to format a system disk or a disk that is currently mounted.

interleave The interleave factor (a decimal number between 1 and 32) to be assigned. The default is 1.

Description:

Detailed information about FORMC formatting is provided in subsequent sections in this chapter. See also the sections in the beginning of this chapter on “General Formatting Information” and “Interleave Factor Calculation”.

Caution The FO command causes an entire tape or disk volume to be formatted. If you specify a disk LU, the entire disk volume is formatted.

Help (?)

Purpose: Displays all the legal FORMC commands.

Syntax: ?

Spare (SP)

Purpose: Spares one or more defective blocks within a track of a CS/80 disk.

Syntax: `SP, media_LU, track`

media_LU The LU of the CS/80 disk or CTD disk cache memory to be spared.

track The track containing bad blocks to be spared.

Description:

You may only use the SP command with CS/80 disks, including the CTD disk cache memory area. The disk cache is divided into four logical tracks; SP can be used to spare one of these cache tracks. Refer to the *RTE-A Driver Reference Manual*, part number 92077-90011, for a description of the disk cache memory scheme.

Sparing using SP is explained in detail in the section “The Sparing Operation” later in this chapter.

Verify (VE)

Purpose: Verifies the integrity of the specified disk or tape.

Syntax: `VE, media_LU [,start, numb]`

media_LU The LU of the CS/80 disk or tape to be verified.

start The number of the disk track or tape block at which verification is to begin.

numb The decimal number of disk tracks or tape blocks to be verified. If you omit the start and numb parameters in the runstring, the entire disk or tape is verified.

Description

Verification is a non-destructive process that checks the integrity of a specified disk LU or CTD tape by examining the checksum information for each block of the LU or tape.

The FORMC Formatting Operation

You must be superuser in order to format a disk or tape, otherwise, the error message

```
INSUFFICIENT CAPABILITY
```

displays, and FORMC aborts.

Entering the LU

When you enter the FO command, FORMC prompts you as follows:

```
/FORMC: CS80 DISK OR TAPE LU?
```

If you are formatting a tape (also called “tape certification” or “tape media initialization”), enter the LU of the CTD drive on which the tape is mounted. You may enter any LU on a disk volume, since the entire volume is formatted by FORMC. If the LU you enter does not correspond to a legal CS/80 disk or tape drive, FORMC displays the following message and prompts you again for the LU number:

```
/FORMC: ILLEGAL CS80 DISK OR TAPE LU
```

FORMC issues a warning and prompts you before a currently mounted disk or the system disks can be formatted.

Disk Formatting

The FO command applies to the entire disk volume; for most CS/80 disks, this means the entire disk unit. Since CS/80 disks do not let you format RTE subchannels as with previous disk drives, you must specifically spare any defective disk blocks found with the SP command described earlier.

After you enter the correct disk LU, you must indicate whether you are using a single-sided or double-sided floppy disk:

```
/FORMC: SINGLE-SIDED FLOPPY (Y/N)?
```

FORMC

Enter “Y” if your disk is single-sided or “N” if your disk is double-sided.

The program then asks if you want FORMC to assign a default record interleave factor:

```
/FORMC: INTERLEAVE FACTOR FOR OPTIMAL THROUGHPUT (Y,N)?
```

Accepting the Default Record Interleave Factor

A “Y” response indicates that you want to use the default interleave factor of 1. This is sufficient for most applications. (Refer also to the earlier section on interleave factor calculation to determine the proper interleave factor for your use.)

Note

If you are using FORMC on double-sided microfloppies, **DO NOT** use the default (enter “N” at the prompt). Drives that are both fixed and microfloppy drives also require an interleave factor of 2. When the next prompt is displayed (see below), enter 2 for the interleave factor.

If you use the default interleave factor of 1, and the timeout is set to 500, timeouts will occur during reads and writes to that LU.

Specifying a Record Interleave Factor

A response of “N” to the prompt above lets you specify a different record interleave factor. When you enter “N”, you are prompted:

```
/FORMC: INTERLEAVE FACTOR?
```

Your response must be a decimal number in the range of 1 to 32. Otherwise, you are prompted again as follows, and FORMC waits for the correct entry:

```
/FORMC: ILLEGAL INTERLEAVE FACTOR, LEGAL RANGE 1 TO 32  
/FORMC: INTERLEAVE FACTOR?
```

Dismounting the LUs

FORMC locates all the disk LUs that are on the same volume and thus, are subject to formatting. If any of the LUs are still mounted, FORMC displays the following messages, then exits:

```
/FORMC: DISMOUNT LU  
<lu 1>, <lu 2> . . . . . , <lu n>  
  
/FORMC: DISMOUNT LU'S AND RESTART  
/FORMC: FINISHED
```

Use the FMGR DC command to dismount the LUs listed in the message and restart FORMC.

The Disk Formatting Process

When the LUs are dismounted, FORMC locks all the LUs to be formatted and displays:

```
/FORMC: DATA WILL BE DESTROYED ON LU  
<lu 1>, <lu 2> . . . . . , <lu n>
```

```
/FORMC: OK TO PROCEED (Y,N)?
```

Respond “Y” to confirm that you want FORMC to format the entire disk volume. FORMC then issues the following message and exits:

```
/FORMC: DISK FORMATTING COMPLETED
```

If you enter N to the confirmation prompt above, FORMC unlocks all LUs and exits without formatting any disks.

Formatting a System Disk

If you are formatting a system disk, FORMC warns you with the following message and waits for your response to the prompt:

```
/FORMC WARNING: SYSTEM WILL BE DESTROYED.  
NOW IS THE TIME TO REPLACE THE SYSTEM CARTRIDGE.
```

```
/FORMC: OK TO PROCEED (Y,N)?
```

If you want formatting to proceed, replace the removable system disk with the disk to be formatted and enter “Y”. Be aware that FORMC formats the entire disk volume, including the system subchannels, if you enter “Y” and do not remove the system disk.

If you enter “N”, FORMC terminates.

If you are using FORMC on 3-1/2 inch, double-sided, flexible disks, you may see the following warning message:

```
/FORMC WARNING: THE DOUBLE-SIDED 3-1/2 INCH FLOPPY THAT YOU JUST  
FORMATTED IS NOT USABLE IN THE INTEGRATED MICRO/1000 FLOPPY DRIVE.  
FORMATTING THIS FLOPPY IN A MICRO/1000 FLOPPY DRIVE, USING FORMC,  
PROVIDES COMPATIBILITY FOR BOTH EXTERNAL AND INTEGRATED FLOPPY  
DRIVES.
```

Refer to the *Small Disk User's Guide*, part number 5958-9152, for more details.

When formatting is complete, FORMC issues the message:

```
/FORMC: REPLACE SYSTEM CARTRIDGE TO RETURN TO SYSTEM.  
/FORMC: OK TO RETURN (Y,N)?
```

Enter “Y” or “N”. After the system disk is replaced, a “Y” response causes a return to the host system; otherwise, FORMC displays the following message and executes an infinite loop:

```
/FORMC: FINISHED
```

When this message displays, you must restore and reboot the system.

FORMC

Tape Formatting

Caution FORMC locks the entire disk when CS/80 cartridge tapes are being certified.

Tape Certifying and Skip Sparing

Formatting a CTD tape is different than formatting a disk. If the tape was not previously certified (the “initialized media” flag on the tape is not set), FORMC performs a tape certification process. This includes writing a known pattern on every block of the tape, then reading it to locate any defective areas. Each 512th block on the tape is allocated as a sparing block, and a table of these blocks is created on the tape. If any defective blocks are discovered, they are flagged and subsequent writes to the tape do not use them. Finally, the initialized media flag is set on the tape.

Skip sparing can take up to twenty minutes for a short tape (DC 150) or more than one hour for a long one (DC 600). RTE-A utilities that access CTD tapes cannot use a tape until the initialized media flag is set.

Re-certification and Spare Conversion

If the tape was certified and the flag is set, you may either re-certify the tape or convert “jump spares” to skip spares. (In jump sparing, a specific defective block is spared to the first spare block following the bad block.) Re-certification is identical to initial formatting, and any data on the tape is destroyed in the process.

The spare conversion process flags the tape blocks that were spared in jump sparing, so they are skipped on the next write to the tape. On the next write, the controller simply renumbers the tape blocks starting at the first good block following a defective block. The spare conversion option results in loss of data for any blocks that were jump spared.

You need not re-certify a tape on an HP 1000 system unless you have serious doubts about the tape’s quality. Defective areas on a tape discovered after certification by a normal read or verification are automatically flagged as defective, and they are skip spared on any subsequent write. Since all HP 1000 utilities use skip sparing, the spare conversion option for HP 1000 CTD tapes is not necessary, though you may use it to reformat other, non-HP 1000 systems that use jump sparing.

Inserting the Cartridge

When you call FO, FORMC determines whether the cartridge is inserted in the drive; this process requires 10 to 15 seconds to execute. If the cartridge is not present, the following messages are displayed, and FORMC exits:

```
/FORMC: READY TAPE AND RESTART  
/FORMC: FINISHED
```

Insert the cartridge into the CTD tape drive and restart the utility.

The Tape Formatting Process

If the cartridge shows an UNINITIALIZED MEDIA status, FORMC issues these messages:

```
/FORMC: TAPE INITIALIZATION(CERTIFICATION) MAY TAKE UP TO AN HOUR  
/FORMC: DATA WILL BE DESTROYED ON TAPE LU xx  
/FORMC: OK TO PROCEED (Y,N)?
```

If you enter “N”, FORMC exits. Otherwise, FORMC proceeds to certify the entire tape, then issues the following message and exits:

```
/FORMC: TAPE FORMATTING COMPLETED
```

If the tape was already initialized, FORMC warns you as follows:

```
/FORMC WARNING: TAPE IS ALREADY INITIALIZED(CERTIFIED)  
/FORMC: DO YOU WANT TO RECERTIFY OR CONVERT SPARES(CE,SP)?
```

When you enter “CE”, FORMC notifies you:

```
/FORMC: TAPE INITIALIZATION(CERTIFICATION) MAY TAKE UP TO AN HOUR  
/FORMC: DATA WILL BE DESTROYED ON TAPE LU xx  
/FORMC: OK TO PROCEED (Y,N)?
```

If you enter “N”, FORMC exits. If you enter “Y”, FORMC re-certifies the tape, issues the following message, and exits:

```
/FORMC: TAPE FORMATTING COMPLETED
```

Converting Spares

If you enter “SP” when FORMC asks if you want to re-certify the tape, FORMC prompts you for confirmation before it proceeds:

```
/FORMC: JUMP SPARES WILL BE CONVERTED TO SKIP SPARES  
/FORMC: DATA MAY BE DESTROYED ON TAPE LU xx  
/FORMC: OK TO PROCEED (Y,N)?
```

If you enter “N”, FORMC exits. Otherwise, FORMC updates the tape spare table by converting jump spares to skip spares, then issues the following message and exits:

```
/FORMC: TAPE FORMATTING COMPLETED
```

In non-interactive mode, FORMC always performs certification on new tapes or re-certification on previously initialized tapes.

FORMC

The FORMC Sparing Operation

You must be superuser in order to perform a spare operation, otherwise, the error message

```
INSUFFICIENT CAPABILITY
```

displays, and FORMC aborts.

Entering the LU

When you call SP, FORMC prompts you:

```
/FORMC: CS80 DISK OR TAPE LU?
```

Enter the LU of the disk that contains defective blocks to be spared. If you enter a CTD tape LU, the disk cache memory area of the CTD is spared. If you enter an illegal LU, FORMC notifies you and prompts again:

```
/FORMC: ILLEGAL CS80 DISK OR TAPE LU
```

If the LU is a tape LU, and the tape does not have a generated disk cache memory, the following message is displayed and FORMC exits:

```
/FORMC: NO CACHE DEFINED FOR LU <number>  
/FORMC: ABORTED
```

When you enter a legal LU number, FORMC prompts you as follows:

```
/FORMC: TRACK TO BE SPARED?
```

Specifying the Track

At the prompt, enter the number of the track that contains the bad block or blocks. If you enter a number outside the legal range, FORMC displays the number of the last track (the number of tracks on the LU minus one) and prompts you again:

```
/FORMC: ILLEGAL TRACK NUMBER, LEGAL RANGE 0 TO <last track>
```

When you enter a legal number, FORMC checks the track for bad blocks. Note that you must use the SP command for each track to be spared.

Sparing Defective Disk Blocks

Sparing removes individual defective disk blocks from active service. Defective blocks on CTD tapes do not require sparing because they are flagged when detected during a normal read operation by the format operation. Those blocks are then automatically spared on any subsequent write operation, a process known as “skip sparing.” However, you may also spare out one or more defective blocks within a track of the CTD disk cache memory area.

Each CS/80 disk has a specified number of physical tracks reserved for use as spares when a defective block is detected. Each physical track also contains a spare block. Spare tracks and blocks are known only to the controller (their existence is transparent to users).

When the program spares a defective block within a physical track, it rearranges the track to bypass the defective block and to use the spare block. If it subsequently spares another defective block on the same physical track, it bypasses the entire track and writes the data to another physical track from the spare track pool.

Once the LU number and track number are known, FORMC performs an error-rate test on the specified track to determine which blocks, if any, are defective on the track. As each defective block is found, it is removed from active service, and the following message is displayed:

```
/FORMC: PHYSICAL BLOCK <address> SPARED FOR <number> BLOCKS
```

The *address* is the physical block address and *number* is the physical number of blocks spared. Since the specified track may cover more than one bad area, the program may issue more than one message. If it must use the next to last spare track, the program issues the following warning:

```
/FORMC WARNING: ONE SPARE PHY. TRACK LEFT. CONTACT SYSTEM MANAGER.
```

Although the program may perform several more sparing operations if the disk only needs one spare block per track, the disk has only one more physical track available for sparing. This warning indicates that there is something seriously wrong with the drive or the disk. A complete diagnostic check of the drive and disk should be performed.

Data Loss During Sparing

If the data in a bad block may be lost when sparing occur, FORMC issues this warning:

```
/FORMC: DATA MAY BE DESTROYED ON LU <number> TRACK  
<track 1> , <track 2> , . . . . . , <track n>
```

```
/FORMC: OK TO PROCEED (Y,N)?
```

If you enter “Y”, sparing proceeds and the data on the tracks indicated by *<track 1>* through *<track n>* is not retained. If you enter “N”, the program terminates.

If the disk LU boundary crosses a physical track that may actually cover another disk LU (such as another FMP directory), FORMC issues the warning:

```
/FORMC: DATA MAY BE DESTROYED ON LU  
<lu 1> , <lu 2> , . . . . . , <lu n>
```

```
/FORMC: OK TO PROCEED (Y,N)?
```

If the data was backed up, you may enter “Y” and have FORMC lock all LUs that might cross physical tracks and perform the sparing operation. If you enter “N”, FORMC terminates.

After a successful sparing operation, FORMC informs you:

```
/FORMC: DISK SPARING COMPLETED
```

If FORMC finds no bad blocks on the specified disk, it issues the following message, then exits:

```
/FORMC: NO BAD BLOCKS - SPARING NOT PERFORMED
```

FORMC

The FORMC Verify Operation

You can use FORMC to verify an entire disk or tape or specific tracks or blocks. A CTD tape block consists of 512 words, and a CS/80 disk block consists of 128 words. Note that all track references are to RTE tracks, not physical disk tracks.

Entering the LU

When you enter the VE command, FORMC prompts:

```
/FORMC: CS80 DISK OR TAPE LU?
```

Enter the LU of the disk or tape to verify. If you enter an illegal number, the program prompts again:

```
/FORMC: ILLEGAL CS80 DISK OR TAPE LU
```

Readying the Tape

Before verifying a CTD tape, FORMC makes sure the cartridge is in the drive. This takes 10 to 15 seconds. If the cartridge is not present, you are informed:

```
/FORMC: READY TAPE AND RESTART  
/FORMC: FINISHED
```

FORMC exits after it verifies the CTD disk cache memory area associated with the tape (if the cache exists). You must insert the cartridge in the drive and restart the utility. If the cartridge was not initialized (formatted), FORMC notifies you and returns to the TASK? prompt:

```
/FORMC: UNINITIALIZED TAPE MEDIA - FORMAT TAPE  
/FORMC: TASK?
```

You must use the FO command to format the tape. Refer to the discussion of the FO command and the section on “The Formatting Operation,” both in this chapter, for more information.

Verifying the Entire Disk LU or Tape

When you insert a formatted cartridge and enter the correct LU, the program prompts you as follows:

```
/FORMC: VERIFY ENTIRE DISK LU (Y,N)?  
or  
/FORMC: VERIFY ENTIRE TAPE (Y,N)?
```

If you enter “Y”, FORMC verifies the entire disk LU or tape; otherwise, FORMC prompts with one of the following messages:

```
/FORMC: START TRACK NUMBER?  
or  
/FORMC: START TAPE BLOCK NUMBER?
```

Specifying the Start Track or Block

If you are not verifying the entire disk LU or tape, enter an integer between zero and the number of tracks or blocks on the medium minus one. For tape verification, the number is between 0 and 65535.

If you enter a negative number, the program uses its complement as the starting block (for example, $-2 = 65534$). If you enter an incorrect number (a number larger than the disk/tape capacity, or in the case of disk verification, a negative number), one of the following messages displays:

```
/FORMC: ILLEGAL TRACK NUMBER, LEGAL RANGE 0 TO <last track>
```

or

```
/FORMC: ILLEGAL BLOCK NUMBER, LEGAL RANGE 0 TO <last block>
```

The *last track* or *last block* is the number of the last track/block minus one.

Specifying the Number of Tracks or Blocks

When you enter the correct number, FORMC displays one of the following prompts:

```
/FORMC: NUMBER OF DISK TRACKS?
```

or

```
/FORMC: NUMBER OF TAPE BLOCKS?
```

You must specify one or more tracks or blocks. If the number you enter is greater than the capacity of the medium, all tracks or blocks following the start number are verified. If you specify a negative number for the number of tape blocks, it is complemented and interpreted as positive (for example, $-5 = 65531$). If you specify zero tracks or blocks, or a negative number of disk tracks, FORMC issues one of the following messages and repeats the prompt:

```
/FORMC: ILLEGAL NUMBER OF TRACKS
```

or

```
/FORMC: ILLEGAL NUMBER OF BLOCKS
```

FORMC

The Verification Process

When you confirm that you want to verify the entire disk LU or tape, or when you enter the correct start track or block and specify the number of tracks and blocks to verify, the program performs the verification.

If you are verifying a CS/80 tape, and it has an associated disk cache memory, the disk cache is also verified.

If FORMC encounters a defective track or block during verification, it notifies you with one of the following messages:

```
/FORMC: LU <disk LU> - BAD TRACK <number>
```

or

```
/FORMC: LU <tape LU> - BAD BLOCK <number>
```

or

```
/FORMC: LU <disk LU> - BAD TRACK <number> IN CACHE
```

where *number* is a decimal integer identifying the defective track or block.

When it completes the verification process, FORMC displays:

```
/FORMC: DISKVERIFY COMPLETED <number> BAD TRACK(S)
```

or

```
/FORMC: TAPE VERIFY COMPLETED <number> BAD BLOCK(S)
```

or

```
/FORMC: TAPE CACHE VERIFY COMPLETED <number> BAD TRACK(S)
```

where *number* is either a decimal integer that defines the number of defective tracks or blocks, or "NO" to indicate that no defective tracks or blocks were encountered.

You can use the SP command, as described in earlier sections of this chapter, to spare any bad tracks after verification.

Loading FORMC

To load FORMC, use the LINK command file #FORMC.

FORMC Error Messages

CS/80 devices return 10 words of status information:

word 1	–	identification field
word 2	–	reject errors (severity #1)
word 3	–	fault errors (severity #2)
word 4	–	access errors (severity #3)
word 5	–	information errors (severity #4)
words 6-10	–	parameter area for words 2-5

One or more error messages may be displayed in the following format:

```
/FORMC: eeeee ERROR xxxxxB
      <mnemonic error code>
IDENTIFICATION FIELD: yyyyyyB QSTAT: q
PARAMETER FIELD P(1) THRU P(10):
ppppppB ppppppB ppppppB ppppppB ppppppB
```

where:

<i>eeeeee</i>	The error type: REJECT, FAULT, ACCESS, INFORMATION.
<i>xxxxxB</i>	The error field, in octal.
<i>yyyyyB</i>	The identification field, in octal.
<i>q</i>	The error-reporting message from the device.
<i>ppppppB</i>	The five-word parameter area, in octal, for the error.

If the program encounters an undefined error bit in one of the status words, the following message replaces the *<mnemonic error code>* in the message:

```
FATAL INTERNAL ERROR
```

Refer to the *RTE-A Driver Reference Manual*, part number 92077-90011, for a description of the error reporting scheme and a complete summary of all mnemonic error codes.

FORMF

FORMF

FORMF is an online and bootable offline maintenance program for floppy disk drives that you can use to verify the integrity of data on a disk unit and to format disks for future use.

FORMF performs the following functions to format a floppy or mini-Winchester disk drive:

- Writes (or rewrites) the track and sector addresses. The disks are divided into physical sectors of 128 words each.
- Interleaves the disk sectors according to the interleave factor (sometimes called fill number), if requested. Interleaving can help reduce access time.
- Identifies bad tracks and makes them invisible. The faulty tracks present no problem unless there are so many that the disk is not usable.

FORMF supports the following disk drives:

HP 9121 S/D	3-1/2 inch microfloppy
HP 9133 A/B/XV	5 and 10 Mbyte Winchester with microfloppy
HP 9134 A/B/XV	5 and 10 Mbyte Winchester
HP 9895	8 inch floppy drive Model 6 built-in 5 inch floppy drive
HP 248x	Integrated 3-1/2 inch single-sided microfloppy
HP 248x	Integrated 3-1/2 inch double-sided microfloppy

Calling FORMF

To call FORMF, enter the following runstring:

```
CI> [RU, ]FORMF[ ,list LU][ , command , disk LU[ , interleave ] ]
```

where:

<i>list LU</i>	The LU of the list device to which FORMF directs messages.
<i>command</i>	The FORMF commands are described below.
<i>disk LU</i>	The LU on the disk unit to be verified or formatted. It is important to note that with some disks, more than one LU can be on one unit. Whichever unit is pointed to by that LU is formatted. The whole unit is formatted, not just the LU.
<i>interleave</i>	The optional interleave factor (a decimal number between 1 and 32) specified with the FO command. The default is an assigned interleave factor calculated for optimal throughput for the drive type being formatted.

If you omit a required parameter or enter one incorrectly, FORMF exits with the messages:

```
/FORMF: ILLEGAL RUNSTRING PARAMETER  
/FORMF: ABORTED
```

If you do not enter any runstring parameters, FORMF runs interactively, prompting you to enter a command as follows:

```
/FORMF: FLOPPY/MINI-WINCHESTER DISK FORMAT UTILITY  
/FORMF: '?' WILL LIST THE LEGAL COMMANDS  
/FORMF: TASK?
```

When you enter a command, FORMF prompts for parameters, which are described below.

Command Execution

Before each command is executed, FORMF issues this message:

```
/FORMF: command executing - <command> <parameters>
```

FORMF terminates after successful command execution or after issuing an appropriate error message when it encounters an error.

Device Driver Status

FORMF tests the disk driver status as it executes commands, and if the driver has timed out, it issues the following messages and exits:

```
/FORMF: DEVICE DRIVER TIMED OUT.  
/FORMF: READY DRIVE ENTER " " RETURN
```

Driver timeout occurs if the disk drive becomes disconnected or if the hardware fails. Refer to the “FORMF Error Messages” section later in this chapter for a description of the device driver status information message.

Break Detection

FORMF checks the break flag during verification (see the discussion of the VE command below) when defective tracks are found. If the flag is set, FORMF exits with these messages:

```
/FORMF: BREAK DETECTED  
/FORMF: ABORTED
```

The program does not check the break flag during the format operation, because formatting must complete without interruption in order to preserve the integrity of the disk. See the discussion of the FO command in this chapter.

FORMF

FORMF Commands

FORMF commands are summarized in Table 2-18 and described in the sections that follow.

Table 2-18. FORMF Commands Summary

Commands	Description
Information Command	
HElp ?	Displays a list of FORMF commands
Configuration Commands	
FO rmat <i>disk lu</i> VE rify <i>disk lu</i>	Formats a disk Verifies the data on the disk
Exit Commands	
AB ort EN d EX it	These three commands terminate program execution

Abort, End, and Exit (AB) (EN) (EX) (/E)

Purpose: Use any of these commands to terminate FORMF.

Syntax: AB, EN, EX, or /E

Description:

You may enter the exit or abort commands in response to any FORMF prompt and cause the utility to terminate execution.

Caution Do not abort FORMF during a format operation, as control information on the disk could be damaged or lost, leaving the disk unusable.

Format (FO)

Purpose: To format the specified disk.

Syntax: `FO , disk LU [, interleave]`

disk LU The LU of the disk to be formatted.

interleave The interleave factor, a decimal number between 1 and 32. If you do not specify an interleave, FORMF assigns an interleave factor that is calculated for optimal throughput for the drive type to be formatted. (See the earlier section on “Interleave Factor Calculation.”)

Description:

Detailed information about FORMF formatting is provided in subsequent sections in this chapter. See also the sections in the beginning of this chapter on “General Formatting Information” and “Interleave Factor Calculation”.

Help (?)

Purpose: Displays a list of FORMF commands.

Syntax: ?

Verify (VE)

Purpose: Verify the integrity of the specified disk.

Syntax: `VE , disk LU`

disk LU The LU of the disk to be verified.

Description:

Detailed information about verification is provided in subsequent sections on “The Verify Operation”.

FORMF

The FORMF Formatting Operation

Before the actual formatting begins, FORMF writes test data and error checking information into the user data and postamble areas as described in the earlier section, “General Formatting Information”. The program tests each block by reading the test data back and calculating the error checking information. If the test passes, the block is normal. If the test fails, the preamble is rewritten to make the block invisible. Each block is tested several times.

FORMF then sends a format code sequence to the disk controller, which performs the actual formatting. In the formatting process, the entire disk unit is formatted, not just the LU specified. Note that all information previously stored on the disk is lost.

If two or more LUs share one unit, part or all of the other LUs are cleared when the unit is formatted. This happens because the controller does not recognize LUs, only the entire unit. This is true for all drives except the multi-volume HP 9133/9134 5Mbyte mini-Winchester disks. Only the formatted unit is cleared; the other units are not affected.

Entering the LU

You may enter any LU on a disk, but remember that the entire unit is formatted, not just the LU specified. If the LU does not correspond to a valid disk drive, FORMF issues an error message and prompts for the disk LU again:

```
/FORMF: ILLEGAL disk LU
```

When you enter a legal disk LU, if FORMF is a system utility and you have superuser capability, the program proceeds; otherwise, the following message displays:

```
/FORMF: INSUFFICIENT CAPABILITY
```

Specifying an Interleave Factor

If the disk LU you specified is a floppy drive, FORMF asks if you want to assign a default interleave factor:

```
/FORMF: INTERLEAVE FACTOR FOR OPTIMAL THROUGHPUT (Y,N)?
```

You may define the order in which logical blocks are located on the disk by specifying an interleave factor, the number of physical blocks to be skipped between logical blocks. With an interleave factor of zero, block n is followed immediately by block $n+1$. With an interleave factor of two, two physical blocks are skipped between block n and block $n+1$. Once the disk is formatted, the interleave is transparent to your software.

Refer to the earlier section on “Interleave Factor Calculation” for a more detailed discussion on calculating the interleave factor.

If you respond by entering “Y”, FORMF supplies an interleave factor calculated for optimal throughput for the disk to be formatted. The default interleave factors are listed during execution.

If you enter “N”, FORMF prompts you to enter an integer between 1 and 32 as follows:

```
/FORMF: INTERLEAVE FACTOR?
```

The Formatting Process

FORMF identifies all the disk LUs that are on the same device as the specified LU and locks them to prevent any interference with the formatting process.

Then FORMF issues the following warning to give you a chance to stop the process before formatting destroys any data that you really want to retain (for example, on the other LUs that share the disk unit):

```
/FORMF: DATA WILL BE DESTROYED ON LU  
<lu 1>, <lu 2>, . . . . . , <lu n>  
/FORMF: DO YOU REALLY WANT TO FORMAT THIS DISK <Y,N>?
```

If you enter “N”, FORMF unlocks the LUs and exits without formatting the disk unit. If you enter “Y”, FORMF sends the command to the controller, and the entire disk unit is formatted. When the formatting procedure is complete, FORMF issues the following message and exits:

```
/FORMF: DISK FORMATTING COMPLETED
```

To format the system disk, a memory based FORMF must be used, since FORMT does not allow formatting of mounted LUs.

The FORMF Verify Operation

Entering the LU

When you enter the VE command, FORMF prompts for the required parameter:

```
/FORMF: DISK LU?
```

Enter the LU of the disk to be verified. If you enter an illegal number, FORMF prompts for the disk LU again:

```
/FORMF: ILLEGAL DISK LU  
/FORMF: DISK LU?
```

Verifying the Disk

After you enter a legal LU, FORMF checks the disk. If a bad track is found, FORMF reports:

```
/FORMF: LU <disk LU> - BAD TRACK <cyl number> <head #>
```

The *cyl number* is the number of the defective track. If no bad tracks are found, the disk is verified, and you are notified upon completion:

```
/FORMF: LU <disk LU> - NO BAD TRACKS FOUND
```

FORMF

FORMF Error Messages

/FORMF: CONTROLLER FAULT

A controller problem has forced FORMF to halt.

/FORMF: DMA PARITY ERROR

A DMA parity error occurred during a transmission to/from the disk.

/FORMF: DRIVE FAULT

A drive fault occurred on the selected drive.

/FORMF: DRIVE NOT CONFIGURED

The drive selected is not configured on the controller board.

/FORMF: FATAL INTERNAL ERROR

An unidentified error has occurred within FORMF. Try running FORMF again.

/FORMF: ILLEGAL ENTRY ENTER “?” FOR HELP!

FORMF is requesting a command; enter ? if you want to display a list of all the FORMF commands.

/FORMF: INSUFFICIENT CAPABILITY

FORMF must be loaded as a system utility and the user must have superuser capability to format a disk.

/FORMF: MAX OF 5 BAD TRACKS EXCEEDED

The flexible disk being formatted has too many bad tracks. Discard the disk.

/FORMF: NO DRIVE ATTACHED

The controller has sensed that no drive is attached to the selected unit.

/FORMF: NO MEDIA IN DRIVE

FORMT has sensed an empty drive.

/FORMF: READY DISK AND ENTER “GO” WHEN READY

The disk drive is not ready, or no disk is in the drive. Insert a disk or make it ready and type a space followed by a carriage return.

/FORMF: REMOVE WRITE PROTECT FROM MEDIA

Remove the write protection to format a flexible disk.

/FORMF: SEEK CHECK DURING OPERATION

The drive has an unrecoverable seek check.

/FORMF: UNABLE TO FIND BUS LU

FORMF requires an HP-IB bus LU to format or verify HP-IB disks.

/FORMF: UNABLE TO FORMAT DISK

FORMF could not recover from an error during a format.

/FORMF: UNABLE TO INITIALIZE DISK

FORMF could not recover from a disk error while initializing a disk track. A preceding error message may give some indication about the fault, or a media problem exists.

/FORMF: UNABLE TO VERIFY AFTER TWELVE PASSES. REPLACE MEDIA OR REPAIR DRIVE.

FORMF was unable to free the media of errors and stabilize it after twelve verify and initialize passes. A media/drive problem exists.

/FORMF: UNABLE TO VERIFY DISK

FORMF could not recover from an error while verifying the disk.

/FORMF: UNFORMATTED MEDIA

You tried to verify unformatted media. Format the media before trying again.

/FORMF: UNSUCCESSFUL LU LOCK

FORMF could not lock all the LUs it needed to maintain bus integrity.

/FORMF: WARNING! POSSIBLE BAD FLOPPY MEDIUM – RETRY FORMF OR DISCARD FLOPPY

An error occurred on the fifth and final format/verify combination. Try to format the disk again. FORMF may identify and eliminate the faulty track. If it fails again, discard the disk.

FORMAT

FORMAT

FORMAT is supplied in relocatable form only. In order to create an offline version of FORMAT and to make sure that no other program accesses a disk LU that is being formatted, you must generate a special system in which there is only one ID segment (the one for FORMAT). This effectively makes FORMAT an offline utility. When FORMAT executes, it actually verifies whether or not there is more than one ID segment. If there is, FORMAT aborts with an error message.

In order to create a FORMAT system under RTE-A, you must relocate %FORMAT and put it into the system via either the BOOTEK or BUILD program. Note that whenever you relocate %FORMAT, you must search the library \$DKLIB and size the program to at least 19 pages.

The configuration instructions are provided in the subsequent section on “Generating a FORMAT System”.

FORMAT provides the following functions:

- Formats a flexible (or floppy) disk. All new flexible disks must be formatted by writing track and sector addresses to them.
- Identifies bad tracks on the disk and marks them. Bad tracks are not spared, but are made invisible to your other programs.
- Initializes an HP 7906/7920/7925 disk LU. This removes the protected status of all the tracks in the LU, cleans up the spare track pool, and prepares the LU for RTE use by sparing any bad tracks. (This is different from the IN command in FMGR.)

Note that this implies you are using an HP disk drive and a disk cartridge that was formatted at the factory. A disk cartridge supplied from another source or direct from assembly will not work.

Formatting is done only on test setups where correct alignment and calibration is maintained.

- Spares individual bad tracks. This process recovers data and assigns a spare track to a device track so that the disk controller can access the track correctly in further operations.
- Verifies a hard or flexible disk. The read and verify process checks the tracks of a flexible or hard disk without altering the data on them.
- Reformats a hard disk. This allows you to clear all accumulated bad tracks by clearing the bad track and spare track information on each track. The disk must have been formatted at least once at the factory before reformatting can be done.

Calling FORMAT

FORMAT runs in interactive mode, prompting you to specify the task to be performed as follows:

```
FORMAT: TASK?
```

When you enter a command, the program prompts you to enter the LU number of the disk to be formatted, initialized, spared, or verified.

FORMAT

Help (??)

Purpose: Prints a list of all valid responses to the current prompt.

Syntax: ??

Initialize an LU (IN)

Purpose: Initializes an entire hard disk LU.

Syntax: IN [*disk lu*]

Description:

The initialization process prepares an entire hard disk LU for use. The previous contents of the disk are overwritten with a pattern of all zeros. Bad tracks are spared to the spare pool at the end of the disk LU.

Refer to the general information sections in the beginning of this chapter to learn more about the differences between initializing and formatting.

See also the subsequent sections under “The FORMAT Initializing Operation” for a more detailed discussion of using the IN command.

Reformat a Hard Disk (RE)

Purpose: Reformats a disk, destroying all the previously written data on the disk.

Syntax: RE [*disk lu*]

Description:

The RE command clears all the status bits on the designated disk, including the spare track pool. It rewrites the preamble, writes zeros in the data area, and rewrites the postamble. It does not verify this data, nor does it spare bad tracks. This lets you clear all accumulated bad track bits, including factory designated bad tracks, and clears the spare and protect bits, as well.

Disk reformatting uses the track map table to determine the starting cylinder, head number, number of tracks, and number of spare tracks for each LU. All tracks in the LU, including the spare track pool, are reformatted. The program reports any disk errors that occur during reformatting.

For more detailed information about using the RE command, see the subsequent sections on “The FORMAT Reformatting Operation”.

Spare a Track (SP)

Purpose: Substitutes a spare track for a defective track and transfers as much data as possible from the bad track to the spare.

Syntax: SP [*track#*]

Description:

Sparing applies to individual bad tracks discovered while the disk is in use. A spare track is substituted for the defective track and as much as possible is copied from the bad track. Offset reads are used in the recovery process, and it is often possible to completely recover the data. In cases where this is not possible, usually only a single block is lost.

General information about sparing bad tracks is presented in sections at the beginning of this chapter. More information about using the SP command is provided in “The FORMT Sparing Operation” later in this chapter.

Verify an LU (VE)

Purpose: Verifies the specified LU and reports defective tracks.

Syntax: VE [*disk lu*]

Description:

Disk verification uses the track map table to determine the starting cylinder, head number, number of tracks, and number of spare tracks for each LU. Tracks in the spare pool are used only as replacements for the defective tracks. They are not verified otherwise.

See also the subsequent sections on “The FORMT Verify Operation”.

Generating a FORMT System

The FORMT utility was designed as a stand-alone utility in order to reduce the possibility of corrupting your disks. This means that you must generate FORMT into a special system whose sole purpose is formatting.

Follow the steps listed below to generate this system under RTE-A:

1. Make a copy of the current system’s generation answer file. Name it FSYS.ANS or something easy to remember or recognize.
2. Reduce the number of ID segments in the system to 1 (one). There are probably 20 to 30 ID segments at present. Change the entry from “ID, 30”, for example, to “ID, 1”.
3. If a MAC disk (HP 7906, 7920, or 7925) is being formatted, follow steps a, b, and c below. Otherwise, skip to step 4.
 - a. Add an entry for %IDM37 in the system relocation phase in the generation answer file.
 - b. Remove the entry for %ID*37 from the driver partition phase. You cannot partition the new MAC disk driver.

FORMAT

- c. Change all references to %ID*37 in the answer file to %IDM37. Module %IDM37 is supplied with the Primary System (part number 92077-16700).
 4. Run the generator, which creates the type 1 system file and the snap file.
 5. Using the snap file from step 4, run the loader to load %FORMAT, searching \$DKLIB. Note that you must size FORMAT to at least 19 pages in order to provide sufficient buffer space. The loader produces a type 6 file named FORMAT. You can use the link command file #FORMAT.
 6. Use the BUILD utility to build the type 6 FORMAT file and the type 1 system file together into your final type 1 format system. When you run BUILD, remember to specify FORMAT as the startup program by using the ST command.
 7. Place the type 1 format system file that BUILD created on a bootable LU and boot it up.
- If your RTE-A FORMAT generation is successful, you should see the “TASK?” prompt.

The FORMAT Formatting Operation

When formatting occurs, bad tracks are identified and made invisible to your software. Bad tracks are not spared; however, unless the number is excessive, they are no longer of concern.

Note that all information previously written to the disk is lost when the disk is formatted.

Entering the LU

After you enter the FO command and specify the LU, the program verifies that the LU is a flexible disk by examining the information defined in the RTE-A driver parameter area. If there is a discrepancy, the program notifies you:

```
INVALID DISK LU
```

You are then prompted again to enter the disk LU.

Confirming Formatting

Since formatting destroys the contents of the specified disk LU, FORMAT asks you to confirm that you want formatting to proceed:

```
DO YOU REALLY WANT TO FORMAT THIS DISK?
```

Respond by typing YES or NO (the first two characters are sufficient).

Sector Interleaving

You are then asked to specify the number of “fill” sectors (the interleave factor):

```
# OF FILL SECTORS?
```

Enter the number of fill sectors desired, between 0 and 28. If you enter zero, sectors are addressed in the same order as they physically appear on the disk. A number other than zero causes consecutive sector addresses to be separated by the number of fill sectors specified. This is known as sector interleaving. For more information, see the section at the beginning of this chapter on “Interleave Factor Calculation.”

Note For best performance, a fill number of 1 is recommended for the system disk or any disk accessed by Hewlett-Packard supplied utilities, including EDITR, the loader, and so on.

The Formatting Process

FORMT makes five passes over the disk to detect any defective tracks. Because each pass may take several minutes, the program displays the following message at the beginning of each pass:

```
FORMAT PASS-xx
```

After formatting the disk, the program reports the number of good tracks:

```
# OF GOOD TRACKS = xxx
```

The program then repeats the request for a task.

FORMAT

The FORMAT Initializing Operation

Remember that you can only initialize a disk cartridge that was previously formatted at the HP factory.

Entering the LU

After you enter the FO command and specify the LU, the program verifies that the LU is a hard disk. Any attempt to initialize a flexible disk or a mini-Winchester causes the following message to display:

```
INVALID DISK LU
```

You are then prompted again to enter the disk LU.

Confirming Initializing

Since initializing overwrites the previous contents of the disk with a pattern of all zeros, the program asks (except when in batch mode) you to confirm that you want initializing to proceed:

```
DATA WILL BE DESTROYED, OK TO PROCEED?
```

When you respond by entering YES, the program performs all the necessary operations to prepare the LU and reports the following information:

```
BAD TRACKS
```

LU xx	LOGICAL CYL	HEAD	UNIT/ADD
BAD TRACK	xxxx	xxxx	x x/x
SPARED TO	xxxx	xxxx	x x/x

```
xx SPARE TRACKS AVAILABLE
```

The program then repeats the request for the task.

The Initializing Process

When initializing occurs, the program cleans up the tracks in the spare pool. The full block (including the preamble and postamble) is read and status bits are examined. If the D bit is set (indicating a defective track), the program proceeds to the next track. A track from the spare pool is not spared. If the D bit is not set, the program rewrites the full track with zeros, including the S, P, and D bits. Verification is performed by reading back the track. If the track is found to be bad during verification, it is flagged as defective.

The program then reads each track in the data portion of the LU. If the D bit is set or the read is unsuccessful, a flag indicating that sparing is needed will be set.

If sparing is needed, the preamble is prepared for a defective track by setting the D bit and the address of the spare track in the preamble of the defective track. The preamble for a spare track is prepared by setting the S bit and the address of the defective track in the preamble of the spare.

The program write initializes the first track, using the preamble set above (if sparing). If this is a good track, the address is the address of the good track (it points to itself). If this is a spare track, the address is that of the defective track (a backward pointer). If this is a defective track, the address is that of the spare.

The data buffer is all zeros if processing the IN command. The same procedure is used to spare an individual track when the SP command is used, except that the data buffer is copied (whenever possible) from the defective track on a block-by-block basis. In recovering user data, an offset read is performed to pick up data that cannot normally be read with the head aligned to the center of the track.

The program verifies the whole track. If verification is successful, the program continues with the next track. If verification is unsuccessful (bad verify status), the program gets the next spare and does the sparing procedure above.

The FORMT Reformatting Operation

Remember that you can only reformat a disk that was previously reformatted at the HP factory, since the servo and timing data are not reformatted.

Before You Use the RE Command

The directory track of the disk LU is also overwritten with zeros when you reformat a disk; therefore, you should take the following steps before using RE:

1. Back up all data from the LU to be reformatted.
2. Dismount the disk LU (use the FMGR DC command).

Entering the LU

After you follow steps 1 and 2 above, enter the RE command and specify the LU. The program checks the disk LU and its type. If you try to reformat a floppy disk LU, the following message appears:

```
INVALID DISK LU
```

You are then prompted again to enter the disk LU.

The Reformatting Process

Since reformatting destroys all data on the disk, FORMT warns you before continuing:

```
DATA WILL BE DESTROYED, OK TO PROCEED?
```

Enter "YES" or "NO". A YES response causes the reformatting to occur; a NO response causes FORMT to terminate.

After You Reformat the Disk

After the successful reformatting of the disk, use the IN command to prepare the disk for File Manager use. IN spares any tracks that are defective, according to the system track map table. Then use the SP command to spare factory designated bad tracks.

FORMT

Follow these steps:

1. Invoke the IN command, which reports and spares all bad tracks.
2. Mount the disk LU (use the FMGR MC command).
3. Restore the data, if desired.
4. Finally, use the SP command to spare factory designated bad tracks, even if the IN command does not report them as bad. Factory designated bad track records are delivered with the disk pack. If your records are misplaced, contact your Hewlett-Packard representative to obtain them from Hewlett-Packard.

The FORMT Sparing Operation

When you invoke the SP command, the program substitutes a spare track for a defective track and transfers as much data as possible from the bad track to the spare.

Entering the Track Number

After you enter the SP command and specify the LU, the program responds:

```
TRACK TO BE SPARED?
```

Enter the track (logical track) found defective on the disk LU. If the track does not lie within the bounds of the disk LU specified, the following message appears:

```
INVALID TRACK #
```

FORMT then prompts you again for the track to be spared. You may enter ?? when prompted, to see the valid track range for the LU.

The Sparing Process

The program copies data, block by block, from the bad track to the spared track. If multiple tries must be made (at each block) with various head offsets, this operation may take several minutes.) If all the information on the track cannot be recovered, FORMT issues the following warning message, then continues:

```
WARNING! ALL INFORMATION ON TRACK NOT SUCCESSFULLY RECOVERED
```

When the program completes execution, the sparing is reported as follows:

```
BAD TRACKS
```

LU <i>xx</i>	LOGICAL	CYL	HEAD	UNIT/ADDR
BAD TRACK		<i>xxx</i>	<i>xxx</i>	<i>x</i> <i>x/x</i>
SPARED TO		<i>xxx</i>	<i>xxx</i>	<i>x</i> <i>x/x</i>

```
xx SPARE TRACKS AVAILABLE
```

If not in batch mode, the program repeats the TASK? prompt.

The FORMAT Verify Operation

FORMAT proceeds by reading each track in the mapped portion of the LU. If the read is unsuccessful, the track is reported as bad.

The directory track of the disk LU is checked; the spare track pool is not checked.

Entering the LU

After you enter the VE command and specify the LU, the program checks the disk LU and its type. If you try to verify a non-disk LU and are running FORMAT interactively, the following message appears:

```
INVALID DISK LU
```

You are then prompted again to enter the disk LU.

The Verify Process

The program reads each track in the mapped portion of the LU, and if the read is unsuccessful, it reports the track as bad. The directory track of the disk LU is checked, but the spare track pool is not checked.

Bad tracks without a valid spare are reported as follows:

```
BAD TRACKS SUBCHANNEL xx

LU xx    LOGICAL CYL    HEAD    UNIT/ADDR
BAD TRACK    xxxx    xxx    xx      x/x
```

You can use the SP command to spare any defective tracks reported here.

If you are in interactive mode, the program returns to the TASK? prompt when the verification completes; otherwise, FORMAT terminates.

FORMAT

FORMAT Error Messages

The following messages may be issued by FORMAT:

CYLINDER COMPARE ERROR
LU *xx* LOGICAL CYL HEAD UNIT/ADDR
TARGET TRACK *xxxx xxxx xx xx/xx*

The disk cannot seek to the target track. Make sure the disk is formatted. The current task aborts, and FORMAT returns to the TASK? prompt.

MORE THAN 1 ID SEGMENT IN SYSTEM.
FORMAT ABORTED

FORMAT was invoked in a system that was generated with more than one ID segment. This is a fatal error; FORMAT aborts.

ENTER FO(RMAT), IN(ITIALIZE), SP(ARE), VE(RIFY), RE(FORMAT), EN(D)

Entry of the first two characters of any of these commands starts that section of the FORMAT program.

INVALID DISK LU
ENTER DISK LU<256

The disk type is wrong for the operation (for example, you are trying to use the IN command on a flexible disk) or the LU has zero tracks, or you entered a non-numeric parameter.

INVALID DISK SPECIFICATIONS *xx*

The disk controller detects an out-of-bounds condition on a cylinder, head, sector, or unit based on an invalid mapping of the disk LU. The current task aborts, and FORMAT returns to the "TASK?" prompt. The disk logical unit number is *xx*.

LU *XX* LOGICAL CYL HEAD UNIT/ADDR
BAD SPARE *xxxx xxxx xx xx/xx*

A defective spare is encountered in the IN or SP task. The numbers are the same as reported in the CYLINDER COMPARE ERROR, above.

MAX OF 20 BAD TRACKS EXCEEDED

This message displays only when you are formatting a flexible disk. The disk cartridge should be discarded.

NOT ENOUGH ROOM FOR TRACK BUFFER

The memory bounds specified for FORMAT at load time did not provide sufficient free memory to serve as a track buffer for the task and disk LU requested. FORMAT terminates.

OUT OF SPARE TRACKS FOR THIS LU

All spare tracks are used up for this LU. FORMAT returns to the TASK? prompt.

FORMAT

READY DISK – ENTER “ ”,CR

The disk drive is not ready, or no disk is in the drive. Ready the disk and type a space followed by a carriage return.

TURN OFF PROTECT OR READ#ONLY SWITCH – ENTER “ ”,CR

The disk protect or read only switch is on. Or, the message may occur if the floppy in the selected drive has the write notch present. Turn off the switch and enter space, followed by return.

TURN ON FORMAT SWITCH – ENTER “ ”,CR

The utility is formatting the disk, but the format switch is off. Turn it on and type a space followed by a carriage return. This message also occurs if a write operation is attempted to a track with its P bit (protected track) set while the Format switch is off.

WARNING! ALL INFORMATION ON TRACK NOT SUCCESSFULLY RECOVERED

Spare only. The track about to be spared was not successfully preserved. FORMAT continues.

WARNING! POSSIBLE BAD FLOPPY MEDIUM – RETRY FORMAT OR DISCARD FLOPPY

Format only. An error occurred on the fifth and final format/verify combination. Restarting the FO process may mark the track defective unless the medium is truly bad.

FST

FST

FST is a high performance, logical (file-by-file) backup utility. It copies files faster than TF and FC because of its streaming capability. It also performs faster than TF and FC on tape units that do not support streaming. FST supports back ups and restores to and from magnetic tape, CS/80 cartridge tape, DDS media, and archive files on disk.

FST reads files from and writes files to both CI volumes and FMGR cartridges. FST does not, however, back up type 0 files. This utility saves files with all the needed extent information so that you can restore them to their original layout, if desired. FST replaces reserved characters in FMGR file names (for example, “.”, “/”, or “@”) with non-reserved characters, and sends a message to the terminal or log device/file when it renames a file.

A virtual memory scheme, using a scratch file, lets you save or restore a virtually unlimited number of files. The only restriction is imposed by the amount of available disk space where the scratch file exists. Any overflow of the scratch file is reported before the files are saved. You may decide where to locate the scratch file.

FST uses a two-pass approach to its backup and restore process. First, you select the files to back up or restore. The selected file names and some other file information are kept in a directory file on the disk. During this first pass, you can add files, remove files, and list file names in the directory file. You may also perform other functions, such as setting the log device/file or changing the selected tape LU.

The second pass begins after you select all the desired files and set all the other backup/restore parameters. During this pass, the files you selected and the directory file itself are transferred between tape and disk. File names and file masks for FST commands conform to FMP standards.

You can use multiple reels for large backups that require more than a single tape. Because a file can cross tape boundaries, files that are larger than an entire tape and multiple files that require more than one tape can be handled.

Each archive created by FST contains a header, an optional comment file, and a directory file, followed by the files saved from disk. A file header immediately precedes each saved file.

You may run FST interactively or programmatically. Parameters in the command string determine the mode. Interactive and programmatic mode are discussed in the following sections.

Calling FST

You may run FST programmatically by entering all the desired commands in the runstring. When FST executes all the commands, or an unrecoverable error occurs, FST exits. The maximum length of the runstring is 256 characters, the limit imposed by CI.

Use a vertical bar (|) to separate commands in the runstring. You must rename files used in the runstring if their names contain a vertical bar.

FST aborts if an error occurs in programmatic mode. If FST ends abnormally, the error is indicated by a “-1” in the first return parameter kept by CI and FMGR.

In the following example, all the files in the working directory are backed up to LU 8 with the Verify option selected:

```
CI> fst ba @|verify|mt 8|go
```

If you do not enter any commands in the FST runstring, FST assumes you are using interactive mode. Commands are entered at the FST> prompt. FST executes each command before prompting for the next one, and continues to prompt for commands until you enter the EXIT command, as follows:

```
CI> fst
FST> <command>
FST> <command>
.
.
.
FST> ex
CI>
```

You may enter just the first two characters of a command, or up to the whole name; for example, EX, EXI, and EXIT all cause FST to exit.

To specify more than one command on a line, separate the commands by a vertical bar (|) as follows:

```
CI> fst
FST> <command> | <command> | ... | <command>
FST>
```

When FST is scheduled, it will look for the start-up command file, FST.RC, in the user’s home directory. If FST cannot find a FST.RC file in the home directory, it will look for /CMDFILES/FST.RC. Every command in the FST.RC file will then be executed, without the normal command echoing, before any runstring or interactive command is processed. Any valid FST command can exist in the FST.RC file. Note that since FST does not allow nested transfer files, it is not possible to transfer to another file from the FST.RC file.

FST

The FST.RC file can be used to set up defaults for FST. For example, the file can be used to override the normal defaults for any of the FST options. The file can also be used to set up a default output device by specifying the MT command in the file.

When FST exits normally, the \$RETURN parameters are set as follows:

\$RETURN1 = Number of errors encountered during the last save/restore pass.

\$RETURN2 = Number of verify errors during the last verify pass.

\$RETURN3 = Number of warnings generated during the last pass.

\$RETURN4 = Number of files saved.

\$RETURN5 = Number of files restored.

FST Commands

Table 2-20 summarizes the FST executable commands.

Table 2-20. FST Commands Summary

Commands	Description
Information Commands	
HE lp (or ?) SH ow	Provides a summary of commands and syntaxes Displays the DF, MT, TI, SC, LL, and option settings
Backup/Restore and Related Commands	
BA ckup <i>mask</i> DF Directory Files <i>file_desc</i> GO Begin Backup/Restore RE store <i>mask</i> SC Select Comment File <i>filename</i> TA UNIX TAR Format TI tle <i>title</i> UN select <i>mask</i>	Selects files to back up Specifies a non-default directory file Begins executing backup/restore Selects files to restore Selects the tape's comment file (backup only) Selects UNIX TAR archive format Specifies a title for the archive (backup only) Unselects files
Listing Commands	
DL Directory List <i>mask</i> LC List Comment File LH List Header LI st Selected Files <i>mask</i> LL Select Log Device/Files <i>device/file</i> LN List Non-Selected Files <i>mask</i>	Displays the archive's directory file Lists the archive's comment file Lists the archive's header Lists the files selected for backup/restore Sets the log device or file for FST activities Lists the non-selected files (restore only)
Tape LU Control Commands	
MT Specify Tape LU or Archive File <i>tape_LU/ filename</i> NE xt <i>append_#</i> PO sition <i>append_#</i> PR evious <i>append_#</i> SD Set Tape Density <i>density</i> SE cure	Sets the magnetic tape LU or archive file Advances the tape to another append Positions the tape to a specific append Rewinds the tape to a previous append Sets the tape density (HP 7974/7978 only) Locks the tape LU and checks the tape status or open archive and checks the file's status
Transfer and Exit Commands	
EX it RU n <i>program</i> TR Transfer to Command File <i>filename</i> / Command Stack	Exits FST Runs an external program Begins executing a transfer file Displays the command stack

FST

Command Stack (/)

Purpose: Displays the FST command stack.

Syntax: / [*n*]

n The optional command line count that specifies the number of command lines from the last command entered to be displayed.

Description:

FST uses the RTE standard command stack, which supports finds, page movement, line marking, and various other operations. For a full description, refer to the online help by typing “? stack” from CI.

The following command stack functionality is available with VC+ only.

FST can optionally save and restore commands to and from a command stack file. If the file FST.STK exists in the user’s home directory, FST restores the command stack from the file at initialization and saves the current command stack to the file when FST terminates.

The \$VISUAL command editing modes that are available in CI can also be used in FST with the CMNDO monitor. To enable this feature, set the exported environment variable \$CMNDO_FST to TRUE. If the variable \$CMNDO_FST is not found, FST looks for the value of the variable \$CMNDO. If neither variable is found, the CMNDO monitor is not used by FST and only the standard RTE command stack is available. For a full description of the \$VISUAL editing modes, refer to the online help by typing “? visual” from CI.

Backup (BA)

Purpose: Selects a file or group of files to back up from disk.

Syntax: BA *mask* [*dest_mask*] [*sec_code*]

mask The file or group of files to be backed up from disk. This can be a disk LU number.

dest_mask The optional mask that changes the characteristics of the files (for example, the path name and file type extension) as they are saved on the archive.

sec_code The system master security code.

Description:

To preserve the security codes when saving FMGR files, specify each file with its individual security code, or enter the system master security code with the mask. If this is not performed, FST will not restore the security code with the FMGR file.

You may execute BA as many times as desired before you start the data transfer to tape. Note that you may not use both the BA and RE (Restore) commands when you select files to transfer.

Directory File (DF)

Purpose: Specifies the name and location of the directory file.

Syntax: DF *file_desc*

file_desc The partial file descriptor that can include a path name, file name, and block size.

Description:

The file descriptor can specify any directory or FMGR cartridge that is not write-protected. FST uses the default path name, file name, or size of the directory file if you do not enter them. The default location of the directory file is the /SCRATCH global directory. If /SCRATCH does not exist, FST creates a directory file on the first available FMGR cartridge. The default file name is a unique file name created by the FmpOpenScratch call. The default size is 500 blocks.

Each selected file requires a minimum of two blocks in the directory file. For large backups that require more than 500 blocks for the directory file, extents are needed for the directory file, and directory file access is slower. To avoid extents, specify the size of the directory file.

Although FST creates a directory file if you do not, DF lets you create directory files for a particular location, name or size. Note that DF can be used only when no files have yet been selected using BA or RE. The following example specifies location, name and size:

```
FST> df /biglu/fstdirfile::::10000
```

The next example just specifies the location:

```
FST> df /lonelydisk/
```

List Directory (DL)

Purpose: Lists the directory of files on the archive.

Syntax: DL [*mask*]

mask The optional mask that specifies a file or group of files in the archive directory. This can be a disk LU number.

Description:

The directory of the archive is displayed and logged to the user log device/file. If you do not specify a mask, the entire directory on the archive is displayed; otherwise, only those files in the archive directory that match the mask are displayed. DL does not modify the list of selected files already obtained by the BA or RE commands.

FST

SHEMA Size (EM)

Purpose: Modifies the SHEMA partition size.

Syntax: EM *no._pages*

no._pages The number of SHEMA pages. This value can range from 50 to 125.

The EM command can only be used as the first command in an FST.RC file. (See the “Calling FST” section for more information about the use of a startup command file.) This command can be used to change the default SHEMA size that is used by the FST and FSTP programs. If this command is not issued, a default value of 125 pages is used for the SHEMA size. Using less than 125 pages of SHEMA may be necessary on a system with memory constraints. See the “Shareable EMA” section on page 2-205 for more information.

Exit (EX)

Purpose: Exits FST.

Syntax: EX

Description:

This command returns you to the environment from which FST was run. See the section on “The Keep (K) Option” later in this chapter for information on exiting without waiting for the tape to rewind.

Begin Backup/Restore (GO)

Purpose: Begins the data transfers to or from the archive.

Syntax: GO

Description:

After you select all the files for the backup/restore, use GO to start the actual transfer of the files.

Help (HE)

Purpose: Displays the command help information.

Syntax: HE [*command*]

 ? [*command*]

command Any FST command or option.

Description:

If you do not specify the optional command, general help information is displayed; otherwise, information for the specified command or option is shown.

List Comment File (LC)

Purpose: Lists the comment file of the archive.

Syntax: LC

Description:

The comment file from the archive is displayed and logged to the log device/file.

List Header (LH)

Purpose: Lists the archive header.

Syntax: LH

Description:

The archive header (format, title, and creation date) is displayed and logged to the log device/file.

List Selected Files (LI)

Purpose: Displays all the files that were selected for backup/restore.

Syntax: LI [*mask*]

mask The optional mask that specifies a file or group of files in the directory file to be displayed. This can be a disk LU number.

Description:

If you do not specify a mask, all selected files in the directory file are displayed and logged to the log device/file. Otherwise, just those files that match the specified mask are displayed and logged. Refer to the “Disk Directory File” section later in this chapter for a description of the directory file.

Select Log Device/File (LL)

Purpose: Changes or selects the log device/file.

Syntax: LL *device* | *file* [a] [o]

device | *file* The device or file to which the output from the listing commands and FST messages is routed.

a The option to append to the specified log file, if one exists.

o The option to overwrite the specified log file, if it exists.

Description:

Using a log file lets you check the file listings and FST messages following a backup/restore. Specifying “LL 1” routes the output to your terminal.

FST

List Non-Selected Files (LN) (Restore Only)

Purpose: Displays the non-selected files; that is, the files that are on the archive but are not being restored to disk.

Syntax: LN [*mask*]

mask The optional mask that specifies a file or group of files on the tape. This can be a disk LU number.

Description:

If you do not specify a mask, all non-selected files in the directory file are displayed and logged to the log device/file. Otherwise, only those that match the mask are displayed and logged. This command should be used only during a restore operation.

Specify Tape LU/Archive File (MT)

Purpose: Selects the tape Logical Unit (LU) number or an archive file name.

Syntax: MT [*tape_LU* | *filename*]

tape_LU The selected tape LU number.

filename The file descriptor of the archive file.

Description:

If you do not specify a tape LU, the current tape LU is set to 0. If you specify an LU other than a tape LU, an error is reported, and the LU value is set to 0. If you specify an LU other than the current tape LU and the current tape LU is locked by FST, FST unlocks the current tape LU and may take it offline before it selects the new LU. Refer to the section on the Keep option later in this chapter for information on when the tape LU is taken offline.

If you specify an illegal tape LU or LU 0, you must use MT to set a legal tape LU before you can execute a backup or restore operation.

If an archive file is specified instead of a tape LU, FST uses this archive file for backups or restores. Archive files are type 1 files and the data can be in either FST format or TAR format. Using an archive file with FST is essentially the same as using a tape. The main difference is that the tape control commands NE, PO, PR and SD are not supported with archive files and the append feature is not supported.

TAR archives are written with a blocking factor of 20 (where one TAR block equals 512 bytes).

FST archive files can be either type 1 or type 2 files. By default, FST creates archive files as type 2 with a record length of 5120 words. Note that prior to Revision 6100, FST supported only type 1 archive files. Therefore, if you are creating an archive file (using Revision 6100) that will be restored using a version of FST prior to Revision 6100, you must specify a type 1 archive file of the appropriate size or copy the type 2 archive file to a type 1 archive file before performing the restore.

Next (NE)

Purpose: Advances the tape to another append of data.

Syntax: NE [*append_#*]

append_# The number of appends to move the tape forward.

Description:

If you do not enter an append number, the tape moves forward to the next append of data. This lets you move the tape and examine various backups that were appended to the tape. If the append number is larger than the number of remaining appends, the tape is positioned at the last append.

Position (PO)

Purpose: Positions the tape at a specific append.

Syntax: PO [*append_#*]

append_# The specific append number to which the tape is moved.

Description:

If you do not specify an append number, the current position is reported. The main backup of the tape is position zero (0). The first append is position one (1), and so on. If you specify an append number larger than the number of appends on the tape, the tape is positioned at the last append.

Previous (PR)

Purpose: Rewinds the tape to a previous append of data.

Syntax: PR [*append_#*]

append_# The number of appends to move the tape backward.

Description:

If you do not specify an append number, the tape is rewound to the previous append of data. This lets you move the tape and examine various backups that were appended to the tape. If you specify an append number that is larger than the number of previous appends on the tape, the tape is positioned at the beginning of the tape.

FST

Restore (RE)

Purpose: Selects a file or group of files from the mounted FST archive to restore to the disk.

Syntax: RE [*mask*] [*dest_mask*] [GR|EG|AG]

mask The optional mask that specifies the file or group of files to be restored to disk. If you do not use a mask, all the files on the archive are restored. This can be a disk LU number.

dest_mask The optional mask that renames the characteristics of the files (for example, the path name and file type extension) as they are restored to disk. This can be a disk LU number that is used by FST if it needs to create a global directory that does not already exist on the system.

GR Begin group restore.

EG End group restore.

AG Abort group restore.

Description:

You may execute RE as many times as desired before you start the data transfer from the archive. Note, however, that RE cannot be used with BA.

A mask is used to select a specific file or group of files. If you do not use a mask, all the files on the archive are restored, with the same path and, if possible, on the original disk LU.

If you are trying to selectively restore many files from a backup tape with a large directory file, the selection process can take quite a long time before the data transfer even begins. This is because when each RE command is executed, a pass is made through the entire directory file. Note that the files are not actually restored until you issue the GO command. Note that with TAR and TF format the GR option is not allowed. RE is handled differently in these cases. See the section TAR or TF Compatibility for more information.

You may use Group commands to speed up large restores. When FST encounters a GR command, it keeps track of all subsequent RE commands, but does not execute them immediately. Instead, they are all executed at once, as a single operation, when the EG (End Group) command is encountered. Thus FST can make a single pass through the directory file, matching all the masks within the group.

Streaming may not occur when you use RE, depending upon the type of tape drive, the size of the SHEMA (shareable EMA) buffer, and the size of the files being restored. However, restore operations imply extraordinary circumstances, such as a corrupt disk or a system that is down, and thus occur much less frequently than backup operations.

Run (RU)

Purpose: Runs a program external to FST.

Syntax: `RU prog_name`

prog_name The name of the program to run.

Description:

RU provides more flexibility when you run FST, as it lets you execute non-FST commands without exiting FST and losing the current FST settings and selections.

For example, you can obtain a directory listing from disk by using the CI DL command with the following command string:

```
FST> ru dl /progs/@.ftn
```

Or, you can create a comment file to save to the tape from within FST by running EDIT/1000 with the following command string:

```
FST> ru edit /my_files/fst/ comments
```

Select Comment File (SC) (Backup Only)

Purpose: Selects the comment file for the archive.

Syntax: `SC [filename]`

filename The file name of the comment file for the archive you are backing up.

Description:

If you do not specify a file name, the current comment file is no longer selected. Comment files are useful for detailed archive identification or restoration instructions.

Set Tape Density (SD) (Backup Only)

Purpose: Sets the tape density for the HP mag tape streaming tape drives.

Syntax: `SD [density]`

density The desired density in bpi for writing to the tape.

Description:

If you do not specify a density number, the current tape drive density setting is displayed. Setting the tape density does not apply to appends; all data on one tape must be at the same density. Refer to your tape drive manual for the proper tape drive density setting.

FST

Secure (SE)

Purpose: Secures (locks) the tape LU and checks the tape status or secures (opens) the archive file and checks the file's status.

Syntax: SE

Description:

SE locks the tape LU specified in the MT command and makes sure the tape is online. SE lets you immediately protect an online tape from other users if it is not write-protected. Using MT to specify a different tape LU unlocks the current tape LU.

Show (SH)

Purpose: Displays the user-selected states of the FST program.

Syntax: SH

Description:

SH displays the states of the option commands, the tape LU or archive file selected, the file count with the number of 128-word blocks and kilobytes represented by the file count, your title and comment file (backups only), the name of the directory file, the log file, and the amount of tape footage needed for the tape access.

TAR (TA)

Purpose: Specifies the UNIX TAR format for the archive.

Syntax: TA, [ON|OFF|A|B][,C]

- | | |
|-----|--|
| ON | Turns ON the TA option. For backups, only type 4 files are converted to a UNIX file format. For restores, all selected files are restored as converted ASCII files. ON is the default. |
| OFF | Turns OFF the TA option. |
| A | (ASCII) Turns ON the TA option and, for backups, converts type 3 and above files to UNIX file format. Restores are handled the same as with TAR ON. |
| B | (Binary) Turns ON the TA option for binary file formats. For backups, files are not converted to UNIX file format and are saved as blocks of data. For restores, all selected files are restored as blocks of data from the archive. |
| C | Allows the selection of case sensitive file names from a TAR archive. |

Description:

Although TA resembles an option, it is a command and thus cannot be entered in a string of options.

When you back up files, you must select all files under the same format. In other words, you may not back up some files in TAR format and other files in FST format on the same archive.

When you restore a TAR archive, FST sets the TA option to ON automatically upon recognizing the archive format. On restores, the default setting restores and converts all files as ASCII files. If a binary restore is required, you must specify the B parameter.

When the C option is enabled, the FST 'RE' command is case sensitive. Also, the FST 'DL' command preserves the case of the file names on the TAR archive when the directory of files on the archive is displayed. Since UNIX files are case sensitive, it is possible to have multiple files on the same archive which, when shifted to uppercase, result in the same name.

For example, to restore both of the "readme" files from a TAR archive that contains 'README' and 'readme':

```
FST> ta,,c
FST> re,README,uppercase
FST> re,readme,lowercase
FST> go
```

```
Tape format:    TAR
```

```
Copying README:::4:1 to UPPERCASE:::4:1
Copying readme:::4:1 to LOWERCASE:::4:1
```

Title (TI) (Backup Only)

Purpose: Specifies a title for the archive header.

Syntax: TI *title*

title The text, up to 72 characters long, that describes the contents of the backup.

Description:

You can use the LH command later to examine the title for the current backup in order to help identify the archive. If you do not specify a title, the last BA command mask entered is used as the default.

Transfer to Command File (TR)

Purpose: Transfers control to a command file.

Syntax: TR *filename*

filename The name of the command file.

Description:

FST executes commands from the file you specify. This is advantageous when you use the same sequence of commands frequently. A command file is quicker to reference, removes the

FST

possibility of typing errors, and does not require operator intervention to execute a series of commands.

To use other commands along with TR in a runstring, specify them before the TR command. FST does not execute any commands in a command string that occur after TR.

Command lines that begin with an asterisk (*) are ignored by FST and can be used to include comments in a command file.

Command files cannot be nested.

Unselect (UN)

Purpose: Removes a file or a group of files from selection in the directory file.

Syntax: UN [*mask*]

mask The optional mask that specifies the file or group of files to unselect from the directory file. (This can be a disk LU number.)

Description:

If you do not specify a mask, UN unselects all the files and purges the directory file from the disk. This command only modifies the directory file and does not affect the files on tape or disk. You can use UN as many times as needed before you actually begin the data transfer to/from the archive.

FST Options

Options enhance the usability of backing up and restoring files. Table 2-21 provides a summary of the FST options. Note that all options do not apply to all commands. You may specify the options, in any order, by entering the first character of the option or up to the entire option name. For example, B, BR, BRI, BRIE, and BRIEF all specify the Brief option.

The SH command displays the current state of the options. All options are initially OFF. Currently selected options are ON.

Set options by entering one or more options on a line and specifying ON or OFF. ON is the default. You may specify an option on the same line as an FST command; to do so, separate the option from the command by a vertical bar (|). You need not set all the options to be used for one backup or restore operation on the same line. The syntax for setting options in interactive mode is as follows:

```
FST> option [option]...[option] [ON|OFF] [option]...[option] [ON|OFF]
```

This allows multiple options to be set ON or OFF in one line.

You may also specify options in the FST runstring as follows:

```
CI> fst option [option]...[option] [ON|OFF] [option]...[option]
```

The examples that follow show how to set options ON or OFF (the examples are entered in interactive mode, but options can also be entered in the FST runstring from the CI> prompt, as shown above):

Example 1: Set Verify option ON.

```
FST> verify on
```

Example 2: Set Append and Clear options ON.

```
FST> a c
```

Example 3: Set Brief, Keep, and Yes options ON; set Lock, Original, and Purge options OFF.

```
FST> b on lock off keep y on orig pu off
```

FST

Table 2-21. FST Command Options Summary

Options	Description
Append	Append this backup to the data already on the tape.
Brief	Only show errors and status messages.
Clear	Clear the disk file's backup bits.
Duplicate	Replace duplicate files.
Faulty	Restores files from a partially overwritten tape.
Inhibit	Inhibit the tape rewind between backups.
Keep	Keep tape online when backup/restore is complete.
Lock	Lock any disk LUs used.
MinDir	Minimize the FST directory file size during restores.
Normal	Backup symbolic links as normal files (follow links).
Original	Restore files to their original main size.
Purge	Purge the disk files after backing up the files.
Quiet	Report messages only to the log device/file.
RwndOff	Rewind and go offline at exit.
SrchApp	Search through appends during RESTORE.
Update	Replace duplicate file if file has been updated.
Verify	Verify the files during the backup/restore.
Whole	Back up all the blocks reserved for the file.
Yes	Write over the tape without asking.
Z	Pause during restore phase for disk full errors.

Append (A) (Backup Only)

Append adds the backed-up files at the end of the tape contents, rather than replacing the previous files. You can only append to tapes that you previously backed up using FST. The Append option is used only with the BA command.

If the mounted tape is not an FST tape, but the Yes option is ON, the tape is overwritten.

Brief (B)

Whenever you do a backup or restore, a “Copying” message is displayed and logged to the log device/file for each file copied. If you specify the Brief option, these messages are not displayed, and only the start and stop of the backup/restore, along with any errors that occur during the process, are shown.

Clear (C)

This option clears the backup bit of copied files that were backed up or restored. In the hierarchical (CI) file system, each file has a backup bit, which indicates whether the file was backed up. When you specify Clear, Verify is automatically ON. If you do not want to verify the backup, Verify may be turned OFF without affecting the status of the Clear option.

Duplicate (D) (Restore Only)

The Duplicate option allows you to replace a file on the disk that has the same name as a file being restored from the tape. If you do not specify the Duplicate option, files with duplicate names are not restored.

When Duplicate is ON, Update is turned OFF, and vice versa.

Note that if you are restoring duplicate files with the Verify option on, FST restores all the duplicate files using temporary file names. The original files are purged and the temporary files are renamed *after* the files are verified. This can require up to twice as much disk space when restoring duplicate files.

Faulty (F) (Restore Only)

The Faulty option lets you restore FST files from a partially overwritten tape, in which the original directory file and probably some actual file data were partly or completely destroyed. Simply position just beyond the overwritten portion of the tape, then build a directory file from the remaining, uncorrupted data. For more detailed information, see the section “Rescuing Files from an Overwritten Tape.”

Inhibit (I)

This option inhibits the normal rewind that occurs after a backup. This does not inhibit the rewind that occurs after FST exits. The Inhibit option is useful when making a backup tape consisting of several appends.

FST

Keep (K)

The Keep option causes the tape to remain online/loaded after you exit FST or select another tape LU. Online refers to 1/2 inch magnetic tapes; loaded refers to CTD (cartridge tape drive) tapes. Note that if the tape is left online/loaded, another user can write over the tape if it is not write-protected. Leaving the tape online also lets FST exit without waiting for the current rewind to complete (except for HP 797x streaming mag tapes).

FST uses the following rules for leaving the tape online/loaded:

- If the tape is write-protected, it is left online/loaded, independent of the Keep option setting.
- If the tape is not write-protected and the Keep option is not set, the tape is taken offline/unloaded.
- If the tape is not write-protected and the Keep option is set, the tape is left online/loaded.

Note that the Keep option assists when using the HP 35401A Autochanger. Refer to the “Multiple Reels” section in this chapter for more information.

Lock (L) (Backup Only)

Setting the Lock option ON locks the disk LUs that are accessed during the backup selection. Setting this option OFF unlocks all locked disk LUs. The Lock option is useful because of the two-phase process that FST uses for its backup operation.

There is a time lapse between the first pass (when all the file information is collected for each file) and the second pass (when the files are transferred). If any file changes during this interval, incorrect or incomplete data may be saved to the tape. The Lock option prevents this by locking the disk LUs needed for collecting data.

This option must be used carefully. When a disk LU is locked, all other users are prevented from accessing that LU. Commonly used disk LUs should only be locked for a short period of time.

MinDir (M) (Restore Only)

The MinDir option is only available when performing an FST restore. Normally, FST restores the entire directory file to disk when restoring any files from an archive. Use of this option limits the user to a single restore mask or group with the RE command.

Normal (N) (Backup Only; VC+ Only)

The Normal option causes FST to back up the data of the file pointed to by a symbolic link. The default behavior of FST is to back up the symbolic link file itself.

Original (O) (Restore Only)

The Original option lets you restore files to the disk with their original main block size. Usually, files of type 3 and above are restored to disk at a block size that contains all the data in the file without creating extents. Sometimes, however, a file must be restored with the same main block size as when it was backed up. You can specify the Original option to accomplish this. This option has no effect on type 1 or type 2 files; they are always restored in their original format.

Note that the Original and Whole (described below) options are not identical. Original determines the main size of a file being restored, while Whole determines how much of a file is backed up.

Purge (P) (Backup Only)

Purge lets you purge the source disk files after they are backed up and verified by FST. When you select Purge, Verify is automatically set; files cannot be purged without verification. The Purge option does not apply to the RE command.

Quiet (Q)

The Quiet option prevents FST output from being displayed on a terminal. Any errors, warnings, or messages are placed in the log device/file. You should specify a log device/file when you use this option; otherwise, it is almost impossible to determine the output of the backup or restore operation. This option is most useful when used programmatically or from a transfer file, but is allowed interactively.

RwndOff (R)

The RwndOff option is mutually exclusive of the Keep option. This option causes the tape to rewind and go offline when FST exits, regardless of the write protect status of the tape.

SrchApp (S) (Restore Only)

The SrchApp option causes FST to search automatically through all of the appends on an FST tape when restoring files. Because this option would typically be used on an FST tape that contains incremental appends, selecting this option automatically enables the UPDATE option. Selecting SrchApp also initializes a group restore (see Restore command). The file names and/or masks from each RE command are stored and used later to search each append's directory file. Restarting grouping, for example,

```
RE [mask] [dest_mask] GR
```

has the effect of reinitializing grouping, and the previous RE commands are invalidated. After the GO command is issued, the directory file for the current append is searched for selected files. After restoring any selected files, the tape is positioned to the next append and the process repeated until the last append is searched. The SrchApp option can only be used with tapes in FST format.

When SrchApp is enabled, the files may not be selectively unselected with the unselect (UN) command. An unselect command (specified without a mask) may be issued to unselect all of the previous selections. This has the same effect as restarting grouping with a RE command.

Update (U)

Update causes FST to restore any duplicate files whose update times on the archive are later than the update times on the disk. This option does not apply to FMGR cartridge restoration, since FMGR files do not have update times.

Update and Duplicate cannot be used at the same time. When the Update option is turned ON, the Duplicate option is turned OFF, and vice versa.

FST

Verify (V)

When you specify the Verify option, FST goes through another pass of the archive after the files are backed up or restored, and compares the data on the tape with the data on the disk to verify that the data was transferred correctly. Streaming may not occur during the verify pass, depending upon the tape drive, the size of the SHEMA buffer, and the size of the files.

Whole (W) (Backup Only)

Normally, FST uses the end-of-file position, specified in the disk directory for the file, to determine how many blocks of data to save to tape. When Whole is ON, FST ignores the end-of-file position and copies all the blocks reserved for the file.

When you back up files to the tape, there is usually no need to save any data beyond the end-of-file position specified in the disk directory entry of each file. When the Whole option is OFF, the end-of-file position in the disk directory is used to calculate how many blocks of data are actually saved to tape. However, sometimes the end-of-file position is corrupt, or does not accurately represent the data to be saved. In that case, the Whole option should be ON.

This option only applies to the BA command. Specify Whole as follows:

1. Set Whole ON.
2. Enter the BA command or commands for the files to be copied with Whole.
3. Set Whole OFF.
4. Enter the BA command for the files to be copied without the Whole option.

Yes (Y)

When you write to an archive that already has data on it, FST asks if you want to write over the archive. You may use the Yes option to suppress this question and have FST copy over the archive automatically.

Z (Z)

The Z option causes FST to pause when the restoration of a file causes a disk full error. This allows the user to free up disk space and restart FST with a "GO" command. FST will then retry the restore of the file that caused the error.

File Masking and Renaming

Although FST masking is designed to be consistent with CI masking, there are some differences, depending upon which command is being executed. Any differences, however, should not affect situations where data could be lost (“unsaved”).

The BA and RE commands refer to copying files, so a D qualifier (described below) is forced into the mask. The DL (List Directory), LI (List Selected Files), LN (List Non-Selected Files), and UN (Unselect) commands simply display file information, so no qualifiers are added by FST. The D, K, N, and S qualifiers are described in the next section.

D, K, N, and S Qualifiers

The D, K, N, and S qualifiers are used when you select files to copy or display. All the qualifiers can be used together; however, the K qualifier overrides the D.

The D qualifier is forced for the BA and RE commands. It has two functions:

- If any directory matches the mask, everything within the directory also matches.
- It preserves the subdirectory path structure of files being copied. (See Example 1 under Backing Up or Example 1 under Restoring in the next sections.)

The N qualifier is almost the reverse of D. N prevents directories from matching and causes subdirectory structures to be “unpreserved” when used with the D qualifier. Since D is sometimes forced, you can use N to help nullify its effects. (See Example 3 under Backing Up or Example 3 under Restoring in the next sections.)

The S qualifier searches down through the entire directory structure. (See Example 3 under Backing Up in the next section.)

The K qualifier searches down through the entire directory structure and preserves the subdirectory path. (See Example 2 under Backing Up or Example 2 under Restoring in the next sections.)

Backing Up

When you select files to save, use the same rules as the CI copy command, CO. Files are matched by name, type extension, or any other characteristic given in the mask, and the D qualifier is forced into the mask. File names and paths can be modified by entering a destination mask with the BA command. Qualifiers on the destination mask are ignored.

Following is a directory tree structure and three examples of file selections for backup. The LI command shows each selected file and how it was renamed. Assume the working directory is A1.dir.

FST

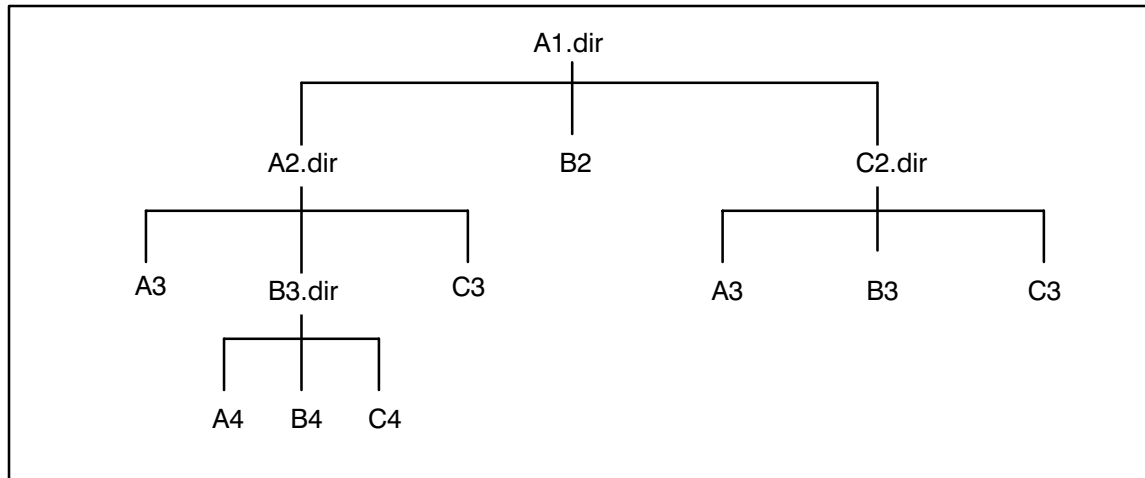


Figure 2-4. Example Directory Tree Structure

Example 1: Back up the files that match the C@ mask.

Note that the files in subdirectory C2.dir are selected even though they do not match C@, because the mask is expanded to C@.@.D and subdirectory structure is preserved under /Z/.

```
FST> ba c@ /z/
4 files selected; 4 total
FST> li

C2.DIR:::2:64:32 to /Z/C2.DIR:::2:64:32
C2/A3:::3:24 to /Z/C2/A3:::3:24
C2/B3:::3:24 to /Z/C2/B3:::3:24
C2/C3:::3:24 to /Z/C2/C3:::3:24

FST>
```

Example 2: Back up the files that match the A@.@.k mask.

Note that the path structure is preserved, every lower path is searched, and the affect of the D qualifier is overridden. That is, files do not match just because their directories do.

```
FST> ba a@.@.k /z/
4 files selected, 4 total
FST> li

A2.DIR:::2:64:32 to /Z/A2.DIR:::2:64:32
A2/A3:::3:24 to /Z/A2/A3:::3:24
A2/B3/A4:::3:24 to /Z/A2/B3/A4:::3:24
C2/A3:::3:24 to /Z/C2/A3:::3:24

FST>
```

Example 3: Back up all the non-directory B files using the masks B@.@.ns.

Note that N causes no directories to match nor subdirectory structures to be preserved, and S causes the search to examine every path.

```
FST> ba b@.@.ns /z/
      3 files selected; 3 total
FST> li

      B2:::3:24 to B2::Z:3:24
      A2/B3/B4:::3:24 to B4::Z:3:24
      C2/B3:::3:24 to B3::Z:3:24

FST>
```

Restoring

When you select files to restore, you follow most of the same rules that apply to the CI command, CO. Files are matched by name, type extension, or any other characteristic given in the mask, and the D qualifier is forced into the mask. You can match file names and paths by entering a destination mask with the RE command. Qualifiers on the destination mask are ignored.

The “mask” parameter in the RE command can also be a disk LU number that selects the files that were backed up from that disk LU. In this case, unless the “dest_mask” parameter is specified as a disk LU number in the RE command, FST tries to restore the file to its original LU. The LU specified in the “dest_mask” parameter is used by FST when FST needs to create a global directory. If the required global directory already exists, FST uses the existing directory.

The differences between restore masking and standard CI masking are as follows:

- Only the D, K, N, and S qualifiers are usable; all other qualifiers are ignored.
- A mask of @, or equivalent, matches every file on the archive, regardless of its path. If you do not enter a mask, the default is @.
- If you do not enter a directory structure with the mask, all paths match and only the file names, type extensions, and other file characteristics are checked for a match (see Example 1 below).
- Global directories, like subdirectories, are preserved as subdirectories if you enter a destination mask (see Example 1 below).

Examples 2, 3, and 4 below provide explanations for the most commonly asked questions about masking, including restoring files to the working directory with path preservation, without path preservation, and with partial path preservation.

Below is a listing of a directory tree structure, backed up to the tape, followed by several examples of file selections for restoring. The LI command displays each file selected and shows how it was renamed. Note that this tree and its files are identical to those shown in the previous section.

FST

```
/A.DIR:::2:64:32
/A1/A2.DIR:::2:64:32
B2::A1:3:24
/A1/C2.DIR:::2:64:32
/A1/A2/A3:::3:24
/A1/A2/B3.DIR:::2:64:32
/A1/A2/C3:::3:24
/A1/A2/B3/A4:::3:24
/A1/A2/B3/B4:::3:24
/A1/A2/B3/C4:::3:24
/A1/C2/A3:::3:24
/A1/C2/B3:::3:24
/A1/C2/C3:::3:24
```

Example 1: Restore the files that match the C@ mask.

Note that the files in subdirectory C2.dir are selected even though they do not match C@, because the mask expands to C@.@.D, and the directory structure is preserved under /Z/.

```
FST> re c@ /z/
6 files selected; 6 total
FST> li

/A1/C2.DIR:::2:64:32 to /Z/A1/C2.DIR:::2:64:32
/A1/A2/C3:::3:24 to /Z/A1/A2/C3:::3:24
/A1/A2/B3/C4:::3:24 to /Z/A1/A2/B3/C4:::3:24
/A1/C2/A3:::3:24 to /Z/A1/C2/A3:::3:24
/A1/C2/B3:::3:24 to /Z/A1/C2/B3:::3:24
/A1/C2/C3:::3:24 to /Z/A1/C2/C3:3:24

FST>
```

Example 2: Restore the files matching the A@.@.k mask, and place them in your working directory.

Note that the path structure is preserved, and the effect of the D qualifier is overridden (that is, files do not match just because their directories do).

```
FST> re a@.@.k
5 files selected; 5 total
FST> li

/A1.DIR:::2:64:32 to A1.DIR :::2:64:32
/A1/A2.DIR:::2:64:32 to A1/A2.DIR:::2:64:32
/A1/A2/A3:::3:24 to A1/A2/A3:::3:24
/A1/A2/B3/A4:::3:24 to A1/A2/B3/A4:::3:24
/A1/C2/A3:::3:24 to A1/C2/A3:::3:24

FST>
```

Example 3: Restore all the non-directory B files using the mask B@.@.ns, and place them into the working directory, flattening the directory structure.

```
FST> re b@.@.ns @
3 files selected; 3 total
FST> li

B2::A1:3:24 to B2:::3:24
/A1/A2/B3/B4:::3:24 to B4:::3:24
/A1/C2/B3:::3:24 to B3:::3:24

FST>
```

Example 4: Restore the files and directory structure under the /A1/A2.dir directory to the working directory.

Note that the A1/A2 directory path is not preserved in the destination.

```
FST> re /a1/a2/@ @
6 files selected; 6 total
FST> li

/A1/A2/A3:::3:24 to A3:::3:24
/A1/A2/B3.DIR:::2:64:32 to B3.DIR:::2:64:32
/A1/A2/C3:::3:24 to C3:::3:24
/A1/A2/B3/A4:::3:24 to B3/A4:::3:24
/A1/A2/B3/B4:::3:24 to B3/B4:::3:24
/A1/A2/B3/C4:::3:24 to B3/C4:::3:24

FST>
```

Incremental Backup

Incremental backup is a procedure that involves periodically backing up all the files and doing frequent backups only of files that were changed since the previous backup. The backup bit in a file's directory entry is used for this, as explained later, so FMGR files do not apply. The initial backup is called a "full backup", and subsequent, selective backups are called "delta backups",

For example, on Friday night, do a full backup of a particular directory. On Monday, Tuesday, Wednesday, and Thursday nights, take delta backups, and append to the same tape that contains the full backup.

Delta backups are done through the use of the backup bit in all hierarchical file directory entries. If the backup bit is set, the file was not backed up; if it is clear, the file was backed up. Whenever a file is created or modified, the backup bit for that file is set. The B qualifier can be used in a mask to select just those files in the working directory whose backup bits are set.

FST

Backup is performed as follows:

1. Use FST to back up all the files, specifying the Clear option to clear all the backup bits of the files that get saved. This is the full backup.
2. Select only those files that have their backup bits set (this is specified by the B qualifier in the file mask), again using the Clear option to clear the backup bits. This is the delta backup.

For example:

```
Saturday: FST> ba /important_data/@
          1000 files selected; 1000 total
          FST> ti IMPORTANT DATA - FULL BACKUP
          FST> C
          Clear ON
          Verify ON
          FST> go
```

```
Monday:  FST> ba /important_data/@.@.b
          90 files selected; 90 total
          FST> ti IMPORTANT DATA - MONDAY DELTA
          FST> c a (the 'a' is optional)
          Append ON
          Clear ON
          Verify ON
          FST> go
```

```
Tuesday: FST> ba /important_data/@.@.b
          120 files selected; 120 total
          FST> ti IMPORTANT DATA - TUESDAY DELTA
          FST> c a (the 'a' is optional)
          Append ON
          Clear ON
          Verify ON
          FST> go
```

Wednesday: etc.

Thursday: etc.

Restoring from Incremental Backups

The SrchApp option can be used to restore selected files from an incremental backup. The SrchApp option automatically moves through each append and searches for files to restore.

Using the previous incremental backup example:

```
FST> s
SrchApp ON
Initialize Group restore
Update ON
FST> v
Verify ON
FST> re @.golf
FST> re four@
FST> re @.day
FST> go
1 files selected; 1 total

Tape format:          FST
Title                 :   IMPORTANT DATA - FULL BACKUP
Created               :   Sat Jun 24, 1989   1:42:14 pm

Copying /IMPORTANT_DATA/NO_GOLF_2.DAY:::4:2:36

Verifying archive

    1 files selected
    1 files restored
    1 files successfully verified.
Positioned at append #1
2 files selected; 2 total

Tape format:          FST
Title                 :   IMPORTANT DATA - MONDAY DELTA
Created               :   Mon Jun 26, 1989   12:57:14 pm

Copying /IMPORTANT_DATA /TROYS_SCORES.GOLF:::4:4:36
Copying /IMPORTANT_DATA /FOUR_EYES.ONLY:::4:1:8

Verifying archive

    2 files selected
    2 files restored
    2 files successfully verified.
Positioned at append #2
1 files selected; 1 total

Tape format:          FST
Title                 :   IMPORTANT DATA - TUESDAY DELTA
Created               :   TUE Jun 27, 1989   1:01:03 pm
Copying /IMPORTANT_DATA/TROYS_SCORES.GOLF:::4:8:36
```

FST

```
Verifying archive
      1 files selected
      1 files restored
      1 files successfully verified.
Positioned at append #0
FST>
```

As an alternative to the `SrchApp` option, you may also choose to use the positioning commands within FST to select the desired append. The FST commands, `NE` (Next), `PR` (Previous), and `PO` (Position), allow the tape movement from one append to another.

The `Update` option restores a file on the tape only if it is newer than the file on the disk. You can use a command file and the `Update` option to devise a general method for restoring just the latest copy of a file.

Using the incremental backup example shown in the previous section, the following general command file example restores a particular file from an incremental backup tape:

```
* Position to the last append (Thursday's)
*
po 4
* Turn on the update and verify options.
Update
Verify
* restore the file (if it's there)
re /important_data/vacation
go
*
* Position to Wednesday's append
*
po 3
re /important_data/vacation
go
*
* Position to Tuesday's append
*
po 2
re /important_data/vacation
go
.
.
.
```

You may also locate the file manually or with a command file, and then position to the correct append and restore the file. The following example shows a command file that is used to find the location of the desired file:

```
d1 /important_data/vacation
ne
d1 /important_data/vacation
ne
```

```

d1 /important_data/vacation
ne
d1 /important_data/vacation
ne
d1 /important_data/vacation

```

When you run this command file, you can see which append has the latest copy of the file, position to that append, and restore the file.

Appending Data

Tape appends are individual backups on the same tape. Each append is separated by tape marks on the tape and has its own tape header, comment file, and directory listing. Each append is independent of the others. You may use the Append option to specify a backup as an append to a tape. The NE, PR, and PO commands are used to examine the individual appends on a tape. The maximum number of appends allowed on a single tape is 1023. Appends are not supported when backing up to an archive file.

Consecutive Backups

The BA command simply appends to the list of already selected files, regardless of the GO command. GO does not clear the selection of files for BA. If you do not want previously selected files for a backup, you must use the UN command to unselect the files. A common mistake is to select a first set of files and copy them to tape using GO, then select a second set, without exiting FST, and copy them to the tape using GO. Thus, the first files are included with the second files on the second copy, because the first files were never unselected.

Multiple Reels

FST supports multiple reels, but handles them differently than other logical backup utilities. When FST backs up files, it splits them across tape boundaries when it reaches the end of a tape. This lets FST back up single files that are too large for one reel.

To restore a single file from multiple reels requires only the reel or reels on which the file is contained. You need not start with the first reel. For example, if you know a file exists on reel 9 of a backup, you can mount the ninth reel and restore the file without having to mount tapes 1 through 8.

The tape format and DL command help determine where a file exists among multiple tapes. The comment file and directory file are at the front of each reel in a multiple reel backup. When the directory file is written to the second reel, the directory is updated, specifying each file that was written to the first reel. Directory updating is also done on all succeeding tapes. This lets you do a DL command on the last tape of a backup to locate a file on any tape in the multiple reel backup.

The Keep option determines the state of the mounted tape when its end is reached during a multiple reel backup. If Keep is ON, the tape remains online/loaded after it is filled. The message “Enter ‘GO’ when a new tape is online/loaded” is then displayed. (Online refers to 1/2 inch magnetic tapes. Loaded refers to CTD and DDS media.)

FST

If Keep is OFF, the tape is taken offline/unloaded after being filled. The message, “FST will continue when the tape is ready (online/loaded)” is then displayed.

You can use the Keep option in several ways. Use Keep OFF when loading and unloading on the HP 35401A Autochanger. Tapes are unloaded when full, and FST continues when the next tape is loaded. If your terminal is not near the tape drives, you can issue GO from the terminal before the tape is online/loaded. If you have set the Yes option ON, you can then go to the tape drive, ready the tape, and the backup will begin. You need not return to your terminal to issue another command.

Tape Loading

If FST is ready to begin the backup or restore operation but the tape unit is not (for example, if the tape unit is offline or the tape is not loaded), FST displays one of two messages, depending upon the state of the Keep option, as follows:

- When the Keep option is OFF, the message “Will continue when tape becomes ready” is displayed. FST begins the backup or restore operation when the tape is loaded correctly.
- When the Keep option is ON, a message that gives the state of the tape unit and tells you to “Type ‘GO’ when ready or ‘BR’ to terminate” is displayed. FST begins the backup or restore operation when you enter one of the commands.

TF Compatibility

FST can restore from TF formatted tapes with FST command functionality; however, FST has less information about TF formatted tapes than about FST tapes. When FST restores files from an FST tape, it uses the directory file on the tape to restore the files. Because a TF formatted tape does not have a directory file, FST cannot determine the files on the tape as easily. Therefore, when FST restores from a TF formatted tape, it stores the file names or masks entered with each RE command into a directory file, which it then uses to search the tape and restore the selected files to disk.

Because FST uses a larger buffer and a faster process than TF, it restores from TF formatted tapes faster than TF, even though FST does not stream when it restores from TF tapes. FST creates a partial directory file of the TF tape during the restore pass, so that tape positioning during the verify pass (on a selective restore) is also faster.

The following example illustrates an FST restore operation from a TF formatted tape:

```
CI> fst
FST> re @.ftn
FST> re makefile
FST> go

Tape format: TF
Title: TapeTitle
Created : Mon Feb 28, 1986 9:40:00 am

Copying file1.ftn
Copying file2.ftn
Copying makefile
FST>
```

TAR Compatibility

FST can read and write archives in TAR format; however, due to the differences between the TAR and FST formats, FST functions in a slightly different manner. Like the TF format, the TAR format does not include a directory file and FST cannot immediately determine the files in an archive. For this reason, when FST restores from a TAR archive, the results of the RE command are not reported after each command. Instead, the file names or masks entered with the RE command are saved and used when the tape is searched for the selected files to be restored.

UNIX Compatibility

Generic RTE/UNIX file system differences are discussed in the UNIX compatibility section in the TF section of this manual.

FST cannot restore ASCII files with record lengths greater than 1024 bytes from TAR archives. Records longer than 1024 bytes will be split into multiple records with a warning issued for each file containing the long records. When creating TAR archives with FST, there is no restriction on the record lengths for type 4 files; however, when a record longer than 1024 bytes is encountered, a warning is issued stating that the record will be split upon restoration by FST. Note that the record structure of the file is still intact and TAR will be able to restore the file without splitting the record.

Files on TAR archives whose names would be illegal RTE file names are renamed by FST. The output of a DL command of a tar archive displays the original file name and the new name created by FST. File names containing reserved characters will have the reserved characters replaced according to Table 2-22.

Table 2-22. Reserved Character Replacements for TAR Archive Files

Reserved Character	Replaced By
<space>	—
.	—
,	—
@	*
—	?
:	
>	^
[(

Note that not all periods are replaced. If the last period in a name results in a legal RTE type extension, it is preserved. File names that begin with numbers are prepended with the underscore character ('_').

To select a renamed file with the RE command, the new file name should be used as the basis for the mask in the RE command. The FST 'DL' command displays the original and new names for any file that is renamed upon restoration.

FST

Rescuing Files from an Overwritten Tape

It is not uncommon to accidentally overwrite a tape. If the overwritten area is smaller than the original backup, you may use the Faulty option, described earlier in this chapter, to restore the data beyond the overwritten area.

Since overwriting a tape destroys the directory file, FST must try to build a new one from the file headers that can still be found. Once a new directory file is built, the operations associated with a normal restore can proceed.

FST assumes that the overwrite ended with at least one EOF mark. After you set the Faulty option ON, position the tape just past the EOF mark that immediately precedes the point where you want FST to start looking for the data that was not overwritten. Then issue the RE command, with a mask if desired. FST builds the new directory file, displays messages about the status of the process, and provides needed information in case the process fails.

The following example shows how to restore whatever files still remain from an FST backup that was partially overwritten by a TF backup, on a magnetic tape.

```
FST> mt 8
FST> f v
Faulty ON
Verify ON
FST> ne 2
Positioned at append #2
FST> re
FAULTY option ON: Assuming tape being restored is partially corrupted.
Searching for a legal tape record
  (sometimes some tape errors are seen here because of parity)
Searching for a valid file header
Scanning tape and building directory file
3 files selected; 3 total
FST> li

f:::4:29:36
g:::3:24:59
h:::4:3:10

FST> go
Copying f:::4:29:36
Copying g:::3:24:59
Copying h:::4:3:10

Verifying tape

FST>
```

It may take a few attempts for you to find the correct file position on the tape, depending upon how many EOF marks actually precede the data you want to restore. If FST reports finding an EOF mark but the process stops, either you did not position the tape to the proper place, or FST did not find any file headers before reaching the next EOF mark.

Note

This process involves unusual tape positioning. Do not try to use the DL command when you restore from overwritten tapes. DL tries to read the directory file from the tape, which may disturb the current positioning.

If the original backup required multiple tapes, the file that crossed the tape boundary cannot be fully restored.

Disk Directory File

When FST backs up to or restores from an archive, it creates a directory file, which contains the names of the files specified in the BA or RE commands along with information about each file that FST needs to perform the backup or restore. Although you cannot display the directory file itself, you may use the LI command to display a list of the files in it.

The size of the directory file on the disk LU limits the size of the backup/restore. You may use the DF command to specify the name and size of the directory file and place the directory file on an LU that is large enough for the current operation. If you do not use DF, FST creates its own directory file when it becomes necessary.

Each time you use the BA command, the contents of the directory file increase. In a restore, because FST uses the directory file on the archive, the entire directory file is copied to disk; therefore, you must specify all the disk space needed for the restore with the first RE command. The directory file on disk is purged when you exit FST.

You cannot move a directory file after it is created. Therefore, use DF before you use BA or RE if the default file descriptor is not adequate. Refer to the earlier discussion of the DF command for directory file default information.

Shareable EMA

FST uses shareable EMA (SHEMA) for its tape buffering. Internally, FST uses between two and five 25-page buffers. The more pages supplied, the faster the speed at which FST backs up files and maintains streaming. However, FST does not use more than five 25-page buffers (125 pages of SHEMA) for tape buffering during streaming.

SHEMA is also used to buffer FMGR file information temporarily when BA is used. If the current SHEMA is not large enough for a FMGR backup selection, more SHEMA is required. The FST EM command can be used to specify the size of SHEMA.

Replacing Reserved Characters

FST replaces reserved characters in FMGR file names with non-reserved characters, and sends a message to the terminal or log device/file that the file was renamed. The reserved characters and their replacements are listed in Table 2-23.

The following exceptions apply when FST replaces reserved characters:

- A period (.) in the middle of a file name is not replaced.
- If there are multiple periods (...) in the middle of a file name, all but the first period in the group are replaced by asterisks.
- If “[” or “>” is the first character in a file name, it is not replaced.

FST

Table 2-23. Reserved Character Replacements for FMGR Files

Reserved Character	Replaced by
.	*
/	!
@	?
[(
>	^

Recommended System Usage

FST uses a directory file to handle the list of all files to back up or restore. Although this speeds up operations (such as doing a DL of the tape or selectively restoring files), the directory file does take up space on the disk and on the tape. Also, the larger the size of the directory file, the longer some operations take to complete.

To avoid slowing down operations and running out of space on the disk LU on which the directory file is located (or to which it is being restored), follow these guidelines when you plan your backup strategy:

- To back up multiple disk LUs that require multiple tapes, divide the operation into saves that fit on a single tape.
- If one disk LU has thousands of files to be saved, save the LU by itself. That is, do not try to back up other LUs with it. Note that the saves can be individual appends on the same tape; therefore, the tape is not being wasted, and you do not need to switch tapes.

Streaming

Streaming is supported on CS/80 cartridge tape drives, DDS tape drives, and the HP 7974/7978 and 7979/7980 Magnetic Tape Drives. Streaming is supported only during backups, not during restores. Restore operations, however, imply extraordinary circumstances, such as a corrupt disk or a system that is down, and occur much less frequently than backup operations.

FST is supported on non-streaming tape drives, but its speed will be much slower than what is available with streaming tape drives. However, FST performs backups and restores on non-streaming tape drives faster than other backup utilities.

Multiple buffers in SHEMA help provide the streaming capabilities during FST backups. Note that files that were written to the buffer may not be copied to the tape before the end of the tape is reached (even if the “Copying ...” message for such a file is displayed or written to the log/device file). In this case, the files are copied to the next tape and the “Copying ...” message for those files is displayed again. Thus, occasionally a file may appear to have been copied to the end of the first tape and to the beginning of the second tape, when actually it was copied only to the second tape.

Streaming is affected by your disk organization. During a backup, if FST accesses many small, scattered files, or files with many scattered extents, continuous streaming is less likely. The larger the files and the fewer extents, the better streaming is maintained.

Streaming is not supported when you back up files in TAR format. Because the process for obtaining the data to be stored to tape is much slower, streaming should not be expected.

FST Format

Each FST archive is written as a single file with 10-Kbyte records. Two file marks are written to mark the end of data on the archive. FST appends are separated by file marks on a tape.

The first 512 bytes of each FST archive is an archive header. An optional comment file can follow the first header. After the first header (and optional comment file) the format of the archive consists entirely of header/data pairs. Each header is 512 bytes and contains information describing the data that follows. The header contains all of the directory information for the file whose data follows it.

The first header/data pair describes the FST directory file. This file contains the headers for all of the files that are included in the archive. When restoring files, FST temporarily restores the directory file as a type 2 file with a 512 byte record length. Following the directory file header/data pair are the header/data pairs for each of the files in the archive. (The header for any file on the archive is saved twice, once in the directory file and again preceding the file's data.) No data is saved when a directory is archived, only a file header indicating the directory's attributes is saved.

Each tape of a multiple tape backup has the comment file and the entire directory file at its head. Each directory maintains the tape number for each file saved on previous tapes. This lets you look at the last tape to determine which tape contains which file.

The format for FST contains all the needed extent information for files. This lets you restore files to the exact needed size, leaving no wasted blocks on the disk, or restore files to their original layout with all the extents or extra space present.

Type 1 and type 2 files are always restored to their original disk format. FMP cannot detect wasted blocks for these types of files or manipulate their extents. For the hierarchical (CI) file system, information such as time stamps and access rights are restored as required.

The FST format and basic header contents are shown in Figure 2-5.

FST

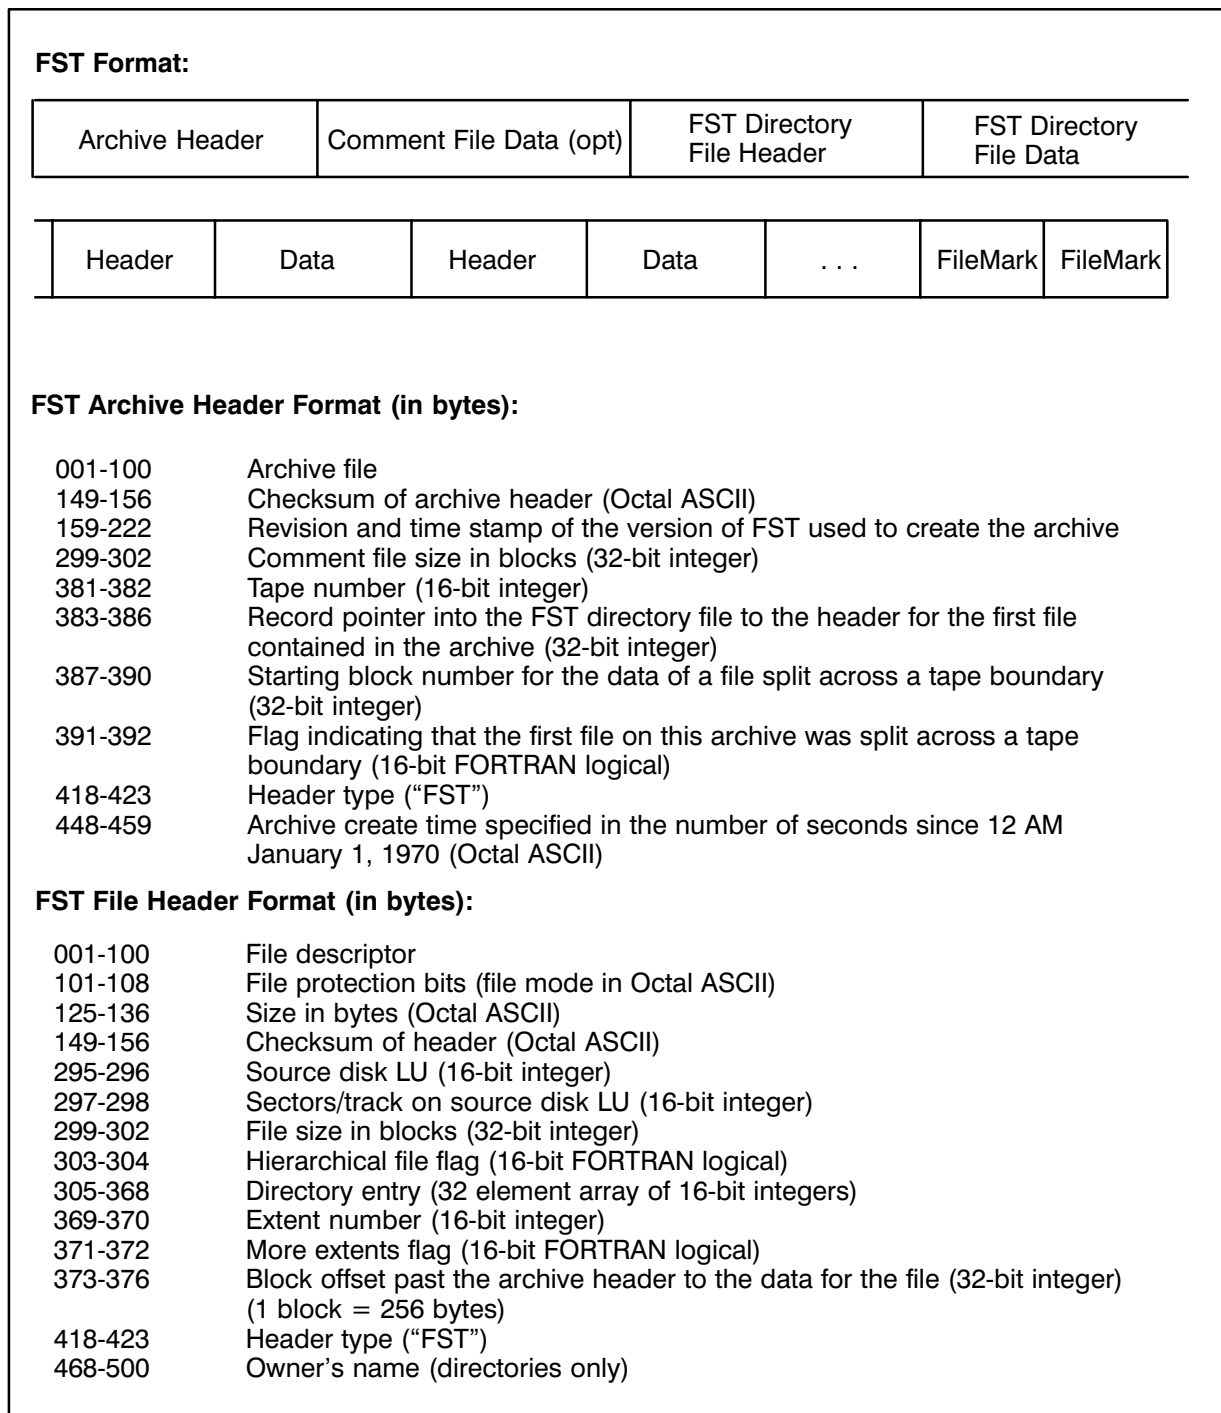


Figure 2-5. FST Format and Header Basics

Installing FST

Two programs, FST and FSTP, are used to facilitate file backup and restore. FST has primary control of all the commands; FSTP handles I/O to and from the archive. For example, when doing a tape backup, FST fills the tape buffers from the disk, while FSTP copies the buffers to tape. In tape backups, FST and FSTP have just enough priority to maintain streaming without preventing other processes from functioning.

To install FST, first link FST and FSTP using the LINK command files FST.LOD and FSTP.LOD. FST generates a unique SHEMA label for each FST/FSTP pair at runtime. A default SHEMA size of 125 pages is used unless specified by the FST EM command in a startup command file.

Place both run files on directory /PROGRAMS or on a FMGR cartridge.

The file >FS000 must be located on the global directory /CATALOGS. If /CATALOGS does not exist, >FS000 can be placed on a FMGR system cartridge.

FST Error Handling

Tape and disk accesses are monitored so that errors can be captured without aborting FST. When you run FST interactively, most errors cause a return to the FST prompt.

The directory file and all specified options remain intact if an error occurs. When you run FST programmatically, errors return a value of -1 in the CI variable \$RETURN1.

If an error causes FST to abort a backup or restore operation and return to the FST> prompt, directory information is not lost. You can reenter the GO command to restart the operation without reselecting the files. We recommend that you use log files for all backups and restores. Errors may occur even when all the specified files are copied.

FST

FST Error Messages and Warnings

The following error messages and warnings may be displayed when you perform a backup or restore using FST:

Aborting FST

FST is being terminated because of a previously reported error.

Appending is not allowed when creating UNIX TAR tapes

The append option was specified for creating a TAR tape, but appends can be done only to FST tapes.

Appends are only allowed on FST tapes

The tape must be in FST format in order to append to it.

Appends are not allowed with archive files

Appends can only be made to FST backups on tape.

Archive file is corrupt: <filename>

While reading the archive file, FST encountered an FMP –12 error and could not continue.

Archive file is not type 1: <filename>

Archive files must be type 1 files. TAR archive files transferred from UNIX machines via FTP must be transferred as binary files.

Break command: process aborted

The current process was aborted because the break flag was set.

Cannot access the log file/device: <filename or LU>

The selected log file/device is not available for use.

Cannot append to the log file: <filename>

The append positioning for the log file was unsuccessful.

Cannot backup: <filename>

A selected file could not be saved to tape successfully.

Cannot backup sparse files across DS: <filename>

You cannot back up sparse files (files with missing extents) across a DS link.

Cannot call a transfer file from a transfer file

The current transfer file contains the FST command to access another transfer file. Transfer files cannot be nested.

Cannot clear backup bit: <filename>

FST could not clear the backup bit of the file backed up or restored.

Cannot create group scratch file. Unable to set SrchApp option

The SrchApp option enables a group restore and 'RE' selections are stored in a scratch file for use during the append search.

Cannot determine the density

The current tape density cannot be determined with the SD command.

Cannot find the directory file on the tape

The directory file could not be located on the tape during tape backup verification.

Cannot find this tape header: <filename>

Incorrect positioning of the tape occurred while searching for a file header. Note that if this error occurs, there is a potential problem with FST.

Cannot lock the tape LU

FST was not able to lock the tape LU.

Cannot create the directory file

The directory file used for file selection could not be created.

Cannot open the archive file: <filename>

The archive file specified in the MT command could not be opened.

Cannot open the selected comment file

FST is not able to open and use the selected comment file.

Cannot open transfer file

FST could not open the specified transfer file.

Cannot position to beginning of data

A rewind to the beginning of an append on a magnetic tape failed. FST did not return to the beginning of the tape append.

Cannot purge: <filename>

FST could not purge the file that was backed up.

Cannot purge the old <filename> in order to rename its replacement

When FST restores a duplicate file to disk, it first copies the file from tape to a scratch file on disk, and gives it a temporary name. If that copy operation is successful, FST purges the original file on disk, and renames the scratch file with the original file name. If FST cannot purge the original file, it cannot rename the scratch file with that file name. In other words, the original file on disk cannot be updated until the old contents can be purged.

FST

Cannot restore: <filename>

FST could not restore a selected file from tape successfully.

Cannot restore <filename> from this tape. It is on tape <tape #> (tape <tape #> is mounted)

The mounted tape belongs to a multi-tape backup, and the selected file exists on a previous tape.

Cannot restore linked files: <filename> linked to <filename>

FST does not restore UNIX hard links or symbolic links from TAR archives.

Cannot restore UNIX device files: <filename>

FST does not restore UNIX device files from TAR archives.

Cannot select file for backup: <filename>

The specified file could not be selected for backup.

Cannot set TAR format: the most recent file selection is from a non-TAR archive. Unselect all files and load a TAR archive for TAR file restoring.

The TAR option cannot be turned on when the currently selected files are from a non-TAR archive.

Cannot successfully rename the restored file <scratch file> to <filename>

When FST restores a duplicate file to disk, it copies the file from tape to a scratch file on disk and gives it a temporary name. If the copy from tape to the scratch file is successful, FST purges the original file on disk, and renames the scratch file with the original file name.

Cannot turn OFF TAR format: the most recent file selection is from a TAR archive. Unselect all masks and load the correct archive for non-TAR use.

The TAR option cannot be turned OFF when the currently selected files are from a TAR archive.

Class I/O between FST and FSTP confused. Current instruction aborted

There is miscommunication between the parent program, FST, and the child program, FSTP, during Class I/O. The current instruction is aborted. If this error occurs, there is a potential problem with FST.

Corrupt comment file on disk

FST cannot obtain the necessary disk information to copy the comment file to the tape.

Corrupt comment file on tape

Corrupt records in the comment file were detected during a listing of the comment file.

Corrupt file. Cannot select: <filename>

The file selected for backup is corrupt.

CTD tape is not initialized

The loaded cartridge tape must be formatted before it can be used.

Directory file failed verify

The directory file failed the verify pass of the backup. This could mean that the directory file became corrupt during the backup operation.

Directory file is corrupt. File selection lost and directory file being purged

The directory file is assumed to be corrupt because the positioning within the directory file failed. The directory file is purged and all previous file selection is lost.

diskerror: Unable to load FST segment

FST was not able to load the needed segment of code.

EMA full. Need to process REstore commands with a GO before continuing.

The number of TF or TAR Restore commands exceeded the EMA limit. Execute the current Restore commands by issuing a GO command; then enter the additional RE commands.

Encountered EOF. Search quitting.

An EOF mark was encountered before any legal file header during an attempt to build a new directory file with the Faulty option.

Erasing the current backup from the tape

The backup in progress is corrupt. A filemark is placed on the tape where the backup began to show that the data following is not valid. The next backup to that tape will begin at that filemark, and the corrupt data will be overwritten.

ERROR: Grouping was never begun.

GR was not entered to begin grouping a set of RE commands, but another grouping command (for example, EG or AG) was entered. You must enter GR before entering other group commands.

Error reading transfer file

FST could not read the next command from the command file successfully.

Error scheduling FSTP:

An FMP error occurred while trying to RP FSTP.

Error scheduling <FSTP rp'd name>: ID segment gone

FSTP was terminated during initialization.

Error scheduling <FSTP rp'd name>: <xxxx> violation

FSTP could not be scheduled successfully.

FST

Error using group scratch file. Cannot process SrchApp option.

The SrchApp option enables a group restore and 'RE' selections are stored in a scratch file for use during the append search.

Extent header missing from archive

FST was expecting the next header on the archive to be for a particular extent, but it was not.

FC tape format: unhandled by FST

The mounted tape is in FC format, which FST does not handle.

File cannot be selected for TAR text conversion: <filename>

The specified file is either inaccessible or corrupt for a TAR ASCII backup.

File failed verify: <filename>

The file on the tape failed to verify (compare identically) with the same file on the disk.

File restored, but not with proper TEXT data: <filename>

The restored TAR file was found to be corrupt on the tape. It was probably a binary file and should not have been restored in ASCII format.

FMP ERROR: <error message>

A report of an error returned from FMP.

FMP ERROR: <error message> – <filename>

A report of an error returned from FMP with the related file name.

Grouping not allowed for non-FST archives

You entered GR to group a set of RE commands when restoring a non-FST archive. Grouping is allowed only for FST archives.

Illegal density

An invalid density was specified with the SD command.

Incorrect usage of command

Incorrect parameters were supplied for an FST command.

Insufficient free space available, size up FST

Free space in the program used for internal buffering is not large enough.

LU <#> is already locked

A mask, with the lock option ON, specified a disk LU that is already locked.

Multiple failures. Search quitting.

While attempting to recover data from an overwritten tape, a good record could not be found.

Need at least 50 pages of EMA to run

Not enough shareable EMA space was linked with FST for proper execution.

No comment file exists on this archive

No comment file exists; therefore, none can be listed with the LC command.

No files selected yet

No files can be backed up, restored, or listed because no successful file selection was completed.

No RNs available: Cannot lock LU <#>

No resource numbers are available to lock the disk LU.

No such archive file: <filename>

The archive file specified in the MT command could not be found.

No tape LU has been selected

A command requiring a specified tape LU was entered, but a legal tape LU was not selected.

Non-FST append found on tape: appending cannot continue

An append to a tape with non-FST data was attempted. FST appends are only allowed on tapes that have all their appends in FST format.

Not a legal tape unit

The MT command was used to specify an illegal tape LU.

Not enough disk space for the archive file: <filename>

The archive file specified in the MT command could not be created in the specified directory. Locate the archive file on a disk with more free space.

Not enough room on this tape to hold the backup. Try another tape.

The mounted tape is too short to hold the current backup directory file and/or comment file.

Not updating: <filename>

The specified file is not being restored because it is older than the disk copy and the Update option is ON.

Option or ON/OFF expected: <unrecognized command>

An illegal word was supplied in an option setting command string.

Owner not set for <directory name>

The original ownership could not be successfully restored to the specified directory.

FST

Protection not set for <filename>

The original read/write protection could not be successfully restored for the specified file.

Setting tape density unsuccessful

The density of the tape unit could not successfully be set to the specified value.

SrchApp mode is available only for FST tapes

The SrchApp mode can be used only when restoring FST tapes.

Tape channel error

The CTD returned a channel error.

TAPE ERROR: <error code>

A tape instruction returned on the no-abort/no-suspend path with the error code in the A- and B-Registers.

Tape FAULT error

The CTD returned a fault error.

Tape headers do not match

The wrong tape is mounted for a multiple tape restore.

Tape is not online

The mounted tape is offline.

Tape LU is down

The tape LU is down.

Tape not ready

The cartridge tape drive unit is not ready for tape access (for example, the tape may be unloaded or positioned to the load point).

Tape sequence wrong: Tape <#> loaded, tape <#> expected

The sequence of tapes mounted during a multiple tape restore operation is incorrect. The mounted tape belongs to the current multiple tape backup, but is out of sequence.

Tape status error

The error bit was set in the A-Register when returning from a tape request call.

Tape write-protected

Files cannot be copied to a write-protected tape.

TAR and non-TAR formats cannot be mixed for backup. Current files must be unselected before TAR format can be used.

The TAR option was turned ON for backup after files were already selected without the TAR option. The formats cannot be mixed.

TAR and non-TAR formats cannot be mixed for backup. Current TAR files selected must be unselected before TAR format can be turned OFF.

An attempt was made to turn the TAR option OFF after files were already selected with the TAR option. The formats cannot be mixed.

TAR selection was used, but the archive is not in TAR format.

The TAR option was turned ON, but the archive is not in TAR format.

The current append is not an FST backup

The tape has been positioned to an append that is not in FST format.

The directory file is already open: cannot assign a new one.

The DF command cannot be used once the directory file has been opened. Note that the UN command can be used to purge the current directory file, but all selected file information must be reselected.

This is not the original tape on which the restore was initiated

A different tape was mounted since the restore operation began. For a restore, FST reads the directory file on the tape during the file selection process. When the GO command is entered, the original tape must still be mounted for a successful restore.

Too many appends specified, tape positioned at last one

There are no more appends on the tape. The tape is still positioned at the last append.

Too many files on this cartridge; size up your EMA

The Extended Memory Area (EMA) is too small for the current file selection for backup from a FMGR cartridge. The EMA must be increased in order to execute it.

Unknown command

The command you entered was not recognized by the program.

Unknown tape format

The mounted tape has a tape format that FST does not recognize (that is, it is not in FST, TF, or TAR format).

Unrecoverable data tape error

The CTD returned an unrecoverable data error.

Unsuccessful path creation

The path of a selected file did not exist and FST could not successfully create the path during a restore operation. The file was not restored.

Warning: <filename> contains record lengths > 1024 words. File saved in TAR mode. (FST will split records during restore.)

A file was saved in TAR format archive that has records with lengths greater than 1024 words. The file was successfully saved in TAR mode; however, a warning is issued because if FST restores this file, the long records will need to be split. (If TAR is used to restore this file, the records will still be intact.)

FST

**Warning: <filename> contains record lengths >1024 words.
FST had to split records during restore of TAR text file.**

While trying to restore a file from a TAR format archive, FST encountered a text file with records greater than 1024 words in length. FST will split records every 1024 words in a record.

**Warning: Illegal FMP name: <illegal filename>
Renamed to : <new filename>**

While reading a TAR format archive, FST encountered a file name that would be an illegal file name. FST will rename the file to a legal file name.

Warning: remainder of command line discarded

Only the commands that precede a TR command, and TR itself, are executed.

Warning: Restore selections cleared

The current append position was changed, so all selected files from there are unselected.

***** Warning *** TAR setting is still on from the restoring. Selected files will be backed up with TAR mode.**

The TAR option is still ON after selecting a TAR archive. Backed up files will be in TAR format.

<xxxx> violation when locking LU <#>

An error occurred when trying to lock a disk LU.

LIF

The LIF utility allows you to interchange files between an HP 1000 system and other HP computer systems. It translates files from the RTE FMP format to a standard Logical Interchange Format (LIF), and vice versa. For example, you can read LIF files and translate them into FMP files that RTE can manipulate. Using LIF, you can write LIF format files on some media and RTE (FMP) files on others.

LIF executes commands that are similar to CI commands. You can store the output to any legal RTE file or LU. LIF is not intended to duplicate the file management capabilities of CI, but to manipulate files on LIF media and interface with FMP. See the sections “Backup Utilities” and “File Interchange on RTE-A” earlier in this chapter to learn which utilities to use for backup and file interchange.

Because LIF conforms fully with the HP standard Logical Interchange Format, the file descriptor (file name and subparameters) used in the standard LIF format differs from RTE’s file descriptor. When media is read from other systems, checks are made to verify the legality of the format. LIF conforms to the proper format when writing and, whenever possible, reads data that does not conform to the standard format.

At times, you may want to write non-LIF file names. This may be useful in preserving naming conventions if both the source and destination systems are capable of handling both conventions. You may partially disable legality checking to allow non-LIF file names to be written. In this case, the utility reports significant differences from the format, but continues to read the media if possible. See the section below on “Naming Conventions”.

Because LIF files are all type 1 (ASCII), all FMP files are copied into type 1 interchange files. Conversely, all files copied from an LIF medium are copied into type 3 FMP files, unless another type is specified in the type subparameter of the file descriptor. Type 1 files need not be ASCII; the type indicates that it is an interchange type file, and thus, has the same record format as other interchange files.

LIF works under RTE-A, Rev. 2326 or later, for all supported HP disks using standard HP disk drivers, including mini-floppies, flexible disks, model 7906 disk drive, CS/80 disk drives (models 7908/7911/7912), and for cartridge tape drives.

Naming Conventions

Not all legal RTE file names are legal LIF file names, because LIF file names consist only of uppercase letters and digits (0-9). When an LIF file name is not specified, the name of the input FMP file is used (the type extension is not included). If that name is not a legal LIF name, legality checking must be suppressed or an error is reported, and the file is not written to the LIF medium.

When files are copied between FMGR (6 character), LIF (10 character), and FMP (16 character) files, excess characters are truncated. Thus, several files with different names under one file system may map into the same file name under another.

If there is a naming conflict, it is treated as a duplicate name error and the file is not copied.

There are three different types of file descriptors, as used in the command parameter description given below. The first type is the FMP file descriptor and is precisely the same as the file descriptor used for file accesses.

LIF

The second type is the LIF file descriptor, which is similar to the first but with a few significant differences. The file name may be up to 10 characters long. It must begin with a capital letter, and the remaining characters must be either capital letters or digits. There is no type extension. LIF accepts any legal FMP file descriptor, including 10 characters, as an LIF file descriptor; however, it may fail to be a legal LIF name (the type extension is deleted). If it is not a legal name, you must disable error checking. It is possible to write illegally named files (refer to the SV command description).

The third type of file descriptor is referred to simply as a namr. This is either an FMP file descriptor or an LIF file descriptor, whichever the context implies. When the file descriptor is of the third type (not specified as either LIF or FMP), the program determines the appropriate type of file descriptor. Ambiguity is resolved by the convention that the mounted LIF medium (there is no more than one at any given time) is searched first. Files that are not clearly specified as either LIF or FMP are searched for (or created) first on the LIF media. If this fails to satisfy the search, the FMP file system is searched using standard FMP calls.

For both types of name, the file descriptor CRN is defined. If that subparameter is a negative number, LIF checks to determine if it matches the mounted LIF media. If so, it is treated as if the reference were to that medium. Otherwise, it is treated as an FMP CRN or global directory name, and handled as a normal FMP file descriptor.

Calling LIF

LIF can be run either interactively or from a transfer file. The entry for interactive operation is as follows:

```
CI> [RU, ] LIF
```

The LIF utility displays the program prompt:

```
LIF:
```

Enter the commands for the desired operation.

To run LIF without any operator intervention, create a command file containing all desired LIF program commands and enter the command file descriptor in the runstring. The runstring is:

```
CI> [RU, ] LIF [ ,command FMP file descriptor [ ,list FMP file descriptor ] ]
```

You may specify a list LU to report error messages or file/device listings. Both the command file and list file must be FMP files because LIF will not run from an LIF file. The default for both files is the display terminal. If the commands are from a file or a non-interactive device, the prompt is suppressed.

LIF Commands

The commands executed by LIF are a subset of the commands executed by CI. In general, the commands work in LIF the same as they do in CI. In some cases, however, the LIF commands do not provide the full range of optional parameters.

Table 2-24 summarizes the LIF commands, which are described in the following sections.

LIF

Copy (CO)

Purpose: Copies all of the files from an LIF medium onto an RTE disk cartridge.

Syntax: *CO* , *-lu* , *destination mask*

lu The list file or device the file is copied to.

destination mask As in the CI CO command. Refer to the *RTE-A User's Manual*, part number 92077-90002, for more information.

Description:

As each file is copied, the name is logged to the list file or device. If errors occur, the related error message follows the file name. All copied files become type 3 FMP files. There is no parameter to specify renaming or copying a subset of the files. You may not copy an FMP disk to an LIF medium because the files on an FMP disk may not be suitable for an LIF medium. Most of the FMP files have illegal names (&, %, ' , ...) and are not of a readily interchangeable type.

LIF file names are truncated to six characters if the destination is a FMGR cartridge. If this results in naming conflicts, files are reported as duplicates and must be renamed.

Directory List (DL)

Purpose: Displays the files that are on the mounted LIF medium and provides some information about them.

Syntax: DL[,*mask*[,*level*]]

mask The wildcard characters that match any letter. Thus if “DL,B----” is entered, LIF lists all the file names with five or fewer letters that start with “B”. The mask can be up to 10 characters long. Note that this mask is not as complex as the mask CI uses.

level The level of directory information requested. The default is level 0, which lists the file names in the directory alphabetically, five names per row. LIF uses the memory space behind the program to store the file names before sorting them, and there may not be enough space to hold all of them. In this case, the level defaults to level 1, which lists only the file names, one per line, in the order in which they appear in the directory. Level 1 may also be legally specified.

Level 2 provides a list that includes the file names in the order they appear in the directory, one per line (only level 0 alphabetizes the names), with their type and size in sectors listed on the same line.

Level 3 includes all of the information available in the directory about the file. After the type and size, the starting track and sector are listed. The date and time stamp (creation time of the file) are listed with the volume number. The volume number indicates which volume of a multi-volume file this medium holds. For a file that spans several media, the volume number indicates which section of the file the current volume holds. If this is the last volume of this file, an asterisk appears before the volume number. (This is normally not used because LIF does not support multi-volume files.) Specifying a level greater than 3 adds a list of purged files to the level 3 information.

For example:

```
LIF: dl,,4
```

NAME VOL NUM	TYPE	# SCTS	TRAK/SEC	DATE/TIME	STAMP
AZUZA12345 *000001	00001	000017	0037 013	13:38:42	10/12/80
BACDEF1492 *000001	00000	000012	0037 011	09:19:26	07/14/81

In this example, file type 00000 identifies the purged file, BACDEF1492.

LIF

Exit (EX)

Purpose: Terminates LIF, closing open files and performing other proper termination functions.

Syntax: EX (or EN, /E, ex, en, /e)

Help (HE)

Purpose: Provides a display of all available LIF commands and their formats.

Syntax: HE (or ?, ??)

Description:

Optional parameters for the LIF commands are specified by brackets. The default values for each command are defined in each related command description.

Initialize (IN)

Purpose: Writes a volume label and blank directory on an LIF medium.

Syntax: IN, *lu* [, *vol label* [, *directory start* [, *directory length*]]]

lu The LIF medium.

vol label The name of the volume, the sector address of the start of the directory, the number of sectors reserved for the directory, and the number of tracks on the medium. The default is "Default".

directory start The default directory start address is 2 (logical sector 2).

directory length The default directory length is 14 sectors.

Description:

You can use IN to prepare a blank medium to use with LIF and to clear the medium of all resident files. IN writes a volume label, an end-of-directory mark over the first file on the medium, and other information not specified by you. When the command is executed, it checks that the LU is an appropriate LIF medium, up and available, and not already mounted to the system. IN then asks:

```
Do You Want To Clear LU xx
```

A Yes answer causes LIF to initialize the medium. Any other response displays the message "Not Initialized" and the LIF: prompt.

Any parameters not specified are replaced by their default values. The default number of tracks is determined from the system tables. After this information is written on the medium, the medium is mounted and made available to LIF. You need not mount an LIF medium before initializing it.

List (LI)

Purpose: LI copies files, either FMP or LIF, to the list file or device.

Syntax: LI ,*filedescriptor*

filedescriptor The file to be copied.

Description:

LI searches the LIF directory first, unless you specified a CRN in the file descriptor. If it does not find the file in the LIF directory, LI searches the RTE directory. When found, the file is listed to the current list file or device. If the file is not found, an error is reported. Note that the LI command does not print line numbers before each line.

Set Logical List Device (LL)

Purpose: LL changes the list file or device.

Syntax: LL[, *FMP filedescriptor*]

FMP filedescriptor The file or device to which listings and error messages are sent.

Description:

If the specified file already exists, it is opened and information is appended to it. If the file does not exist, it is created as default type 3, size 24.

Outputs to the list device are generated by the DL, LI, and HE commands, and other routines (error messages, and so on.) The default list device is the scheduling LU, normally the terminal.

Mount Cartridge (MC)

Purpose: Mounts an LIF cartridge, specifying that it is the cartridge to be referenced by subsequent commands.

Syntax: MC ,*lu*

lu The LIF medium.

Description:

The MC command does not mount an FMP cartridge; it checks that the LU being accessed is an appropriate LIF medium, is up and available, and is not mounted to FMP. Then it reads the volume label from the medium and saves the information that specifies the size and type of the medium, the number of directory tracks, and so on. The medium is also checked for conformity to the LIF standard, and to see if severity is positive. Any errors in format are reported.

The mounted LIF medium is used by all subsequent commands that implicitly refer to the LIF medium. Only one LIF medium can be mounted at a time. If one is already mounted, it is dismounted when another medium is mounted.

LIF

Pack Cartridge (PK)

Purpose: PK packs all files on the LIF medium to recover free space.

Syntax: PK

Description:

The PK command moves the files on the LIF medium next to each other, eliminating the space that was previously occupied by purged files. It also consolidates directory entries, removing the entries for the purged files. Before a PK operation, the purged files are still listed in the directory as type 0 files (shown by entering “DL , , 4”). The PK command removes these entries by sliding all subsequent directory entries down. The end-of-directory mark is also moved down to the last active file.

Purge (PU)

Purpose: PU removes a file from the LIF medium.

Syntax: PU , *LIF filedescriptor*

LIF filedescriptor The name of the file to purge.

Description:

PU changes the type of a file to 0, which in LIF is a purged file. This is the only action the PU command takes. A purged file cannot be listed; it appears only on a level 4 or higher DL and is physically written over when the media is packed.

Rename (RN)

Purpose: RN changes the name of an LIF file to a new one.

Syntax: RN , *old LIF filedescriptor* , *new LIF filedescriptor*

old LIF file descriptor The old file name.

new LIF file descriptor The new file name.

Description:

The RN command rewrites the directory entry by inserting the new name for the file. No other information changes. The new file name is first checked, to make sure it is a legal LIF name, unless the SV level is negative (refer to the description of the SV command for details).

Store (ST)

Purpose: ST creates files from existing files on FMP/LIF medium to FMP/LIF medium.

Syntax: *ST* , *source file descriptor* [, *destination file descriptor*]

<i>destination file descriptor</i>	The CRN specified by you; otherwise, the file is placed on the LIF medium (because it is at the top of the pseudo cartridge list).
--	--

Description:

This command is similar to the FMGR ST command. You may also specify LU numbers. If you specify a CRN as the source file descriptor, it is searched for the file. If CRN is not specified, the LIF medium is searched. If the file is not found, the FMP file system is searched. If the file is still not found, the utility displays:

```
LIF Error: No such file xx
FMP Error: No such file xx
```

If you specify a destination file that already exists, ST is not executed, and an error is reported (duplicate filename). The same naming restrictions apply as in the CO command.

If you do not specify a destination, when the source file is found, the destination file is given the same name but is placed on the medium of the other format. For example, if the source file is found on an FMP disk, that file is stored on the LIF medium; if it is found on the LIF medium, it is copied to the first FMP disk.

Note that the above example does not imply that stores are done only from LIF to FMP and back. FMP files (or LUs) can also be stored to other FMP files (or LUs). It is also legal for you to store LIF files to other LIF files. Thus, all four combinations are possible with this command:

FMP to FMP, FMP to LIF, LIF to FMP, LIF to LIF.

When you specify an FMP file descriptor, the full file descriptor is used, either as a source or a destination. Thus, you can individually specify the type (default is 3), the size (default is twice the number of sectors occupied on the LIF medium), the CRN (default is the working directory or top of the cartridge list), and the security code (default is 0).

The ST command does not let you specify specialized control, format and file-skipping functions, as in the FMGR ST command. Therefore, ST is not normally used for storing FMP files to different types of FMP files. For example, storing a binary relocatable from cartridge tape into a type 5 file does not create a loadable type 5 file.

LIF

Severity (SV)

Purpose: SV modifies the amount of error checking done by the utility. LIF is used to write and read interchange media, checking the media for conformity to the LIF standard. SV allows you to specify how much error checking is required to ensure conformity.

Syntax: SV[, *severity level*]

severity level The severity level can be either positive or negative. Two positive numbers, 0 and 1, can be entered to ensure conformity to LIF standard. These numbers correspond to the standard LIF revision levels, 0 and 1. A negative severity level turns off checking entirely, allowing illegal names and unusual configurations (such as an unusually placed directory).

Description:

Because SV can be used to store FMP files with names that are illegal in LIF, if the destination system can read non-legal names, it may be easier to maintain the illegal but descriptive FMP names. The utility can read any medium it writes. However, other systems may not be able to read files written with a negative severity level. If no severity level is entered, the LIF utility displays its present value:

```
SV Level is S000001
```

Transfer Control (TR)

Purpose: TR lets you transfer control from one RTE file or LU to another. It lets you specify the source of the commands LIF executes.

Syntax: TR[, *FMP filedescriptor*]

FMP file descriptor The command file or LU where the LIF commands are to be entered. If commands are not to come from an interactive device, the prompts are suppressed. The default command file is the scheduling LU.

Description:

There is no “:” before each line. Note that there is no command for calculations (as in the FMGR TR command) or to transfer back, except through an explicit TR command to return to the original file. Thus, there is no command file stacking.

LIF Error Handling

There are numerous error checks throughout this utility. When an error is detected, a message is written to the list device indicating the type of error. All LIF errors are matched as closely as possible to an error number from the FMP error codes. The last error reported is in the RMPAR word 1 on program completion.

The error format is:

LIF: *message*

FMP: *message*

For example:

LIF: No Such File XYZ::-41

FMP: Duplicate filename ABC:123:DB:3

FMP errors are standard errors described in the *RTE-A Programmer's Reference Manual*, part number 92077-90007. Refer to that manual or the *RTE-A User's Manual*, part number 92077-90002, for error codes and definitions.

TF

TF

TF allows you to copy files between disk and tape media, either disk-to-tape or tape-to-disk. You can copy the files to/from a 1600-bpi magnetic tape, CS/80 cartridge tape, or DDS media.

TF is similar to the FC utility, but works with all files, whereas FC works only with FMGR files. TF supports a subset of the FC commands, which provide only tape-related operations. TF does not do disk-to-disk copying. (The CI CO command, described in the *RTE-A User's Manual*, can be used for copying files on disk.)

You can select options for special copy functions, including verifying the copy, replacing duplicate files, suppressing the display of the file names being copied, and appending files to a previously written TF tape.

When you select verification, files are verified by a direct comparison of the disk and tape data after the files are copied. Even when verification is not selected, tape checksums are generated and checked.

Incremental backup makes backing up large numbers of files less time consuming. With incremental backup, you back up all the files occasionally, and just the files that were modified since the last backup more frequently. The modified files are usually appended to the same tape used on the previous backup.

The TF tape format is compatible with the TAR format, used by the UNIX TAR (Tape Archive) utility. TAR is the standard tape format for HP implementations of UNIX. The terms UNIX TAR and TF TAR refer to the original format as written by the UNIX TAR utility and TF's adaptation of that format, respectively. See the discussion of UNIX compatibility in the "Copy Command Options" section in this chapter for more information.

Although TF writes tapes only in TF TAR tape format, it can read tapes in FC format (with some limitations), UNIX TAR, or TF TAR format. You do not need to specify the tape format in the command.

You may run TF interactively or from a command file. Both modes are discussed in the following sections.

Be sure to read the section on "Installing TF" that appears later in this chapter before running TF for the first time.

Calling TF

When you run TF interactively, commands are entered at the TF: prompt. TF executes each command before prompting for the next one, continuing until you enter the EXIT command. For example:

```
CI> tf
tf: <command>
tf: <command>
.
.
.
tf: ex
CI>
```

You may enter just the first two characters, or up to the whole name of a command; for example, EX, EXI, and EXIT all cause TF to exit. Uppercase and lowercase are acceptable.

When you enter a command in the TF runstring, as:

```
CI> tf <command>
```

TF executes the command and terminates.

To run TF with a command file, use the TR command, as:

```
CI> tf tr <command file>
```

<command file> can be a device LU or a disk file descriptor. TF executes each of the commands in the command file, then terminates.

TF ignores command lines that begin with an asterisk (*); you can use them to include comments in a command file.

The syntax conventions defined in the preface of this manual apply to TF, except when file descriptors contain DS locations. In those cases, the “[]” and “>” characters are significant in the entry. Refer to the “Copy Examples” sections in this chapter for more information.

TF Commands

TF uses command sets that let you configure the copy operation, specify names for the tape header and comment files, selectively copy files, group copy commands, transfer to and return from command files, and direct listings to a disk file or list device. Table 2-25 summarizes the commands, which are described in the following sections.

Table 2-25. TF Commands Summary

Commands		Description
Information Command		
Help or ?		Provides a summary of commands and command syntaxes
Copy and Related Commands		
CO Copy Files	<i>filename</i>	Copies files as specified by parameters
DE fault	<i>file_desc</i>	Sets default source, destination, and options for subsequent COPY commands
GR oup	<i>commands</i>	Used for grouping more than one copy command
EG End Group		Ends Group into a single COPY operation
AG Abort Group		Aborts Group
TI tle	<i>title</i>	Establishes title to be used in tape header file
Listing Commands		
LH List Header		Lists header file from TF tape
LL Set List Device	<i>device file</i>	Sets list device or file for LH and DL or File
DL Directory List	<i>mask</i>	Compiles directory list of TF tape
Transfer and Exit Commands		
TR ansfer	<i>filename</i>	Transfers to/return from TF command file
EX it		Exits TF
Comment Command		
*		Identifies following string as comment line

Copy (CO)

Purpose: CO and its related grouping and default commands are used for disk-to-tape (backup) or tape-to-disk (restore) copy operations. The terms “backup” and “disk-to-tape copy” are synonymous in this section, as are the terms “restore” and “tape-to-disk copy”.

Syntax: CO *source destination options*

The parameters are discussed in the sections that follow.

Description:

When files are copied, extents are always combined, making a larger main file. Unused space after the end-of-file is removed, except when copying FMGR files.

TF cannot be used to back up “sparse” files (files with missing extents), such as the VMA backing store file. Any attempt to back up a sparse file results in FMP error –104.

TF cannot be used to back up type 12 byte stream files. The FST program should be used for that purpose.

TF only backs up temporary files if they are explicitly specified for backup. When using a mask to specify the files to be backed up, temporary files are skipped.

CO Command Source and Destination Parameters

Source and destination parameters take several specific forms (described below), depending upon the purpose of the command. Examples are provided in the sections on “Copy Examples” later in this chapter.

- Normal backup, or disk-to-tape copy

A file mask, or multiple masks, selects the files to be backed up. Refer to the “File Masks” section of the *RTE-A User's Manual* for more information. When multiple masks are specified, the list must be enclosed in braces {}. The command's format is:

```
co mask tape [options]
or
co {mask1 mask2 ... maskn} tape [options]
```

Note that if more than one mask selects a file, duplicate copies of the file are written to the tape.

- Normal restore, or tape-to-disk copy

Restore all files from tape. The destination parameter is omitted, which causes files to be restored with the same names and to the same directories from which they were backed up. The command's format is:

```
co tape [ , ,options]
```

- Selective restore

Restore only certain files from the tape specified by the mask or series of masks. Again, the destination parameter is omitted. See the section on “Copy Examples” in this chapter. The command format is:

```
co tape {mask} [ , ,options]
or
co tape {mask1 mask2 ...} [ , ,options]
```

Note that unlike backup, if more than one mask selects a file on a restore operation, the file is only copied once. If more than one copy command in a group has a mask matching a given file, only the first matching copy command has any effect, and determines the destination and options used for that file.

TF

- Restore to different directories or files

You can restore files to directories other than those from which they were saved, or give files different names or type extensions when they are restored. Use the destination mask to specify the directory, name, type extension, security code (FMGR files only), and file size for the files copied from tape. You may specify a mask or list of masks to select specific files or sets of files from the tape, as in an ordinary selective restore. For example:

```
co tape destination_mask [options]
or
co tape {mask} destination_mask [options]
or
co tape {mask1 mask2 . . .} destination_mask [options]
```

- Special purpose backup

A disk-to-tape copy can require that the files on tape are given different file descriptors than they had on disk. You may use either a single source mask or a list of masks, as in a normal backup command. File descriptors are assigned to the files on tape based on the destination mask. The format is:

```
co mask tape {destination_mask} [options]
or
co {mask1 mask2 . . .} tape {destination_mask} [options]
```

Note that a list of destination masks is not allowed. The section on “Grouping Copy Commands” describes how to specify a different destination mask for each source mask.

CO Command Options

The CO command options parameter lets you specify any one or a combination of the options shown below. Options may be specified in any order. When the A, I, V, or Y option is used with GR, the option applies to the entire group if it is selected for any command in the group. Once specified, the K option remains in effect until TF terminates. All other options apply only to the command for which the option was specified. Table 2-26 summarizes the CO command options.

Table 2-26. TF CO Command Options Summary

Option	Description
Append	Adds files to end of current tape
Brief	Suppresses logging of each file name as it is copied
Clear	Clears the backup bit for each file copied to tape
Duplicate	Replaces existing files with files restored from tape
Ignore	Causes errors in the tape header to be ignored
Keep	Prevents TF from taking the tape unit offline
N (Inhibit)	Stops text file conversion from UNIX to FMP
Update	Replaces duplicate files if the existing file has an earlier update time
Verify	Compares files on disk and tape after they are copied
X (UNIX)	Converts type 4 files to UNIX text file format when files are copied to tape
Yes	Suppresses questions about overwriting and appending to the current tape

Append to Tape (A)

Append adds the files being copied to the end of the current tape, rather than replacing the previous contents. Files can be added only to TF tapes; you may not append to FC tapes or UNIX TAR tapes.

The time required to append to magnetic tapes is a function of the number of files already on the tape. Because a forward file search is not required for CTD (the end-of-data position is recorded in the tape header), appending is very efficient with CTD tapes. Refer to the section “Tape Protection” later in this chapter, and the discussion of the K option below, for special precautions to take when you append to tapes.

Brief Logging Mode (B)

This option suppresses the logging of each file name as it is copied, thus reducing the output to the terminal to a small number of messages, plus any error messages.

Clear Backup Bit (C)

The Clear option clears the backup bit for each file copied to tape, thus marking the file as backed up. This is useful when you use incremental backup, in which files are only backed up if they changed since the previous backup (the backup bit is set each time a file is accessed). Refer to the section on “Incremental Backup” in this chapter for more information on how to use the Clear option.

TF

Replace Duplicate Files (D)

The Duplicate option causes existing files to be replaced with files restored from tape if the file names are the same. When you do not use D, a “File already exists” error occurs if you try to restore a file that has the same file name as an existing file, and the file is not restored. Refer to the section “Replacing Duplicate Files” later in this chapter for more information.

Ignore Errors and File Marks (I)

This option is useful in recovering as much data as possible from a partially corrupted (damaged or overwritten) tape. Normally, TF does not continue reading a tape if it does not recognize the header format (“unknown tape format” error), or if a “tape i/o error” occurs when reading the header. The I option causes errors in the tape header to be ignored. TF continues reading the tape even if the header is not valid and ignores any “unexpected eof mark” errors.

This option is valid only for TF and UNIX TAR tapes; it has no effect when reading FC tapes.

Keep Tape online (K)

The Keep option prevents TF from taking the tape unit offline. While this option is convenient, it makes it possible for the tape to be overwritten by another program, because the tape unit remains unlocked and online. Refer to the section on “Tape Protection” in this chapter for more information.

Inhibit UNIX to FMP Text File Conversion During Restore (N)

This option is normally used when you restore binary files from UNIX TAR tapes, but can be used whenever you do not want line-feed characters converted to record boundaries. Unless you specify the N option, this conversion is done for all files restored from UNIX TAR tapes and for type 4 files on TF tapes if you specified the X option when you backed up the files. Refer to the discussion of the X option below for more information.

Update (U)

Update is similar to Duplicate, except that when you use the Update option, duplicate files are only replaced if the update time of the file from the tape is later than the update time of the existing file on disk. See the section on “Replacing Duplicate Files” in this chapter for more information.

Verify Files Copied (V)

Verify compares the files on tape and on disk after they are copied, for both disk-to-tape and tape-to-disk copy commands. The comparison is done after the entire copy is completed. TF uses a scratch file when it copies files to tape with Verify selected. The scratch file is created in the global directory /SCRATCH, if possible. If this directory does not exist, you must create the scratch file in the working directory, or, if there is no working directory, on a FMGR cartridge. Refer to the section on “Installing TF” later in this chapter for more information.

UNIX Compatibility (X)

This option causes type 4 files to be converted to the UNIX text file format (in which line-feeds are used to separate records) when files are copied to tape. No UNIX compatibility option is necessary when copying files from tape to disk. Refer to the section on “UNIX Compatibility” later in this chapter for more information.

Yes (Y)

Yes suppresses the questions, “Do you want to write over this tape (Y/N)?” and “Do you want to append to this tape (Y/N)?”. TF assumes a Y response when this option is specified.

Copy Examples without Subdirectories

The examples in this section are useful if you do not use subdirectories. All the files in the examples are files in global directories or on FMGR cartridges. Two-character names are used for FMGR cartridges (for example, ::X1) and longer names are used for global directories (for example, /PROJECT1). The tape LU is LU 8.

These examples parallel the discussion in the earlier section, “CO Command Source and Destination Parameters.”

- Normal Backup Examples

Example 1: Back up all files in global directory /PROJECT1, with the Verify (V) option selected.

```
co @.@::project1 8 V
```

You may prefer the following notation, which has the same effect:

```
co /project1/@.@ 8 V
```

When the name and type extension are both @, and no mask qualifiers are used, you may drop the @.@ to save typing and abbreviate the commands as:

```
co ::project1 8 V
```

or

```
co /project1/ 8 V
```

All four commands shown above are identical in effect. In the remaining examples, the “::” notation is used. The abbreviated form (with @.@ dropped) is used whenever possible. If a mask qualifier is used, you cannot drop @.@. For example, you cannot abbreviate @.@.X, where the X mask qualifier is used with @.@ to select files with extents (note that ..X is not allowed).

Example 2: Back up all files on cartridge X1.

```
co ::X1 8
```

Example 3: Back up all FORTRAN source files (files with type extension FTN) in global directory PROJECT1, with the B (Brief) option selected.

```
co @.ftn::project1 8 b
```

Example 4: Back up all files on cartridge X1 whose names begin with ‘&’.

```
co &@::X1 8
```

TF

Example 5: Back up all type 4 files in global directory PROJECT1.

```
co ::project1:4 8
```

Example 6: Back up all files in global directory PROJECT1 not accessed since May 1983.

```
co @.a-8305::project1 8
```

Here the A mask qualifier is used to select files with particular access times (year 83, month 05). You may do this to back up old files before purging them to save disk space.

Example 7: Back up all files in the working directory.

```
co @.@ 8
```

If there is no working directory, this command backs up all files on all cartridges in the cartridge list.

Example 8: Back up two files.

```
co {file1 file2} 8
```

Example 9: Back up all FORTRAN, Pascal and MACRO source files in the directory PROJECT1, with the X (UNIX compatibility) and V (Verify) options selected.

```
co {@.ftn::project1 @.pas::project1 @.mac::project1} 8 X V
```

Alternatively, you may use DE to specify the common source directory, as follows:

```
de ::project1  
co {@.ftn @.pas @.mac} 8 X V
```

Note that while DE sets a default source directory, this is NOT the same as setting a default working directory using the CI WD command.

Example 10: Back up files in the global directories PROJECT1, PROJECT2, and SYSTEM.

```
co {::project1 ::project2 ::system}8
```


Example 11: Back up all documentation files (type extension .doc) in global directories PROJECT1, PROJECT2, and PROJECT3 and verify.

```
co {@.doc::project1 @.doc::project2 @.doc::project3}8 V
```

Because the documentation files all have the type extension .DOC, you may use the DE command to save typing as follows:

```
de @.doc
co {::project1 ::project2 ::project3} 8 V
```

You may want to specify more masks than can be typed on a single command line. TF allows 150 characters, but practical considerations usually dictate a maximum of 80 characters. In this case, you can use the GR and EG commands to combine the copy commands into one operation and enter them as follows:

```
de @.doc
gr
co::project1 8 V
co::project2 8 V
co::project3 8 V
eg
```

Because all the directories in the above grouped commands are to be backed up to LU 8 and verified, you may expand the DE command to include the other common properties as follows:

```
de @.doc 8 V
gr
co ::project1
co ::project2
co ::project3
eg
```

Similarly, you may enter the following commands to back up a large number of files. Note that commas are used as placeholders for the omitted source parameter.

```
de,,8 V
gr
co file1
co file2
.
.
.
co lastfile
eg
```

Because DE is so versatile, you can use it to greatly simplify command entry, particularly with GR and EG to copy a large number of files. In the next example, DE is specified several times:

```
gr
de .ftn::dir1 8V
co file1
```

TF

```
co file2.mac
.
.
.
co file50
de .ftn::dir2 8 V
co file51
.
.
.
co file100
eg
de
```

The first DE command specifies that the default file type extension is .FTN, the default source directory is ::DIR1, the default destination is LU 8, and the default option is Verify. Because one of the files, FILE2.MAC, is not a FORTRAN file, that file name is explicitly entered in the CO command, which overrides the default type extension for that file only.

The second DE command changes the default directory to ::DIR2. Because each subsequent DE command cancels the previous one, you must enter the entire set of desired default parameters each time you specify DE.

DE is entered a third time, this time with all null parameters, to cancel the defaults set by the second DE command. This is necessary only if you no longer desire defaults. Defaults remain in force until you override them by entering another DE command, or until TF exits.

Refer also to the discussions of the Default and Group Copy Commands later in this chapter.

- Normal Restore Examples

Example 1: Restore all files from tape LU 8.

```
co 8
```

If you omit the destination parameter, all files are restored to the same directories from which they were backed up (and with the same names and properties).

If the required directories do not exist when you do the restore, TF creates them automatically. You must mount any FMGR cartridges required prior to the restore.

If you do not specify a destination when you restore files, they are restored to the same directory from which they were backed up. Sometimes a directory is not specified when files are backed up, as in the following command:

```
co @.@ 8
```

In this case, the working directory is assumed. If there is no working directory, the FMGR cartridge list is searched for file1.

If a working directory exists, the files are restored to that directory even if it is different from the one in effect when the files were backed up.

Example 2: Restore all the files from tape LU 8 with the Verify option selected.

```
co 8,,V
```

You must separate the source and options parameters with two commas when you do not specify a destination parameter.

- Selective Restore Examples

Example 1: Restore files from tape LU 8, selecting only the FORTRAN source files.

```
co 8{@.ftn}
```

In this example, all files are again restored to their original directories.

Example 2: Restore files from tape LU 8, selecting only those files from global directory PROJECT1.

```
co 8{::project1}
```

The files are restored to the same directory.

Example 3: Restore all FORTRAN, Pascal, and Macro source files, and verify.

```
co 8{@.ftn @.pas @.mac},,V
```

TF

- Restoring to Different Directories or Files

Example 1: Restore all files on LU 8 to the global directory ALLPROJECTS.

```
co 8 ::allprojects
```

All the files are restored to ALLPROJECTS regardless of the directory from which they were backed up.

Example 2: Restore all files on LU 8 to FMGR cartridge X2.

```
co 8 ::X2
```

All file names are truncated to six characters, the maximum allowed under FMGR. Because this is not a function of TF but of the file system, no message is issued if a file name is truncated.

Example 3: Restore all files on LU 8 to the current working directory.

```
co 8 @. @
```

If there is no working directory, the above command restores each file to the first available FMGR cartridge on the cartridge list.

Example 4: Restore all files backed up from global directory PROJECT1 to the global directory PROJECT_XYZ.

```
co 8{::project1} ::project_xyz
```

Example 5: Restore all documentation files backed up from global directories PROJECT1 and PROJECT2 to global directory PROJ_1_AND_2_DOC.

```
co 8{@.doc::project1 @.doc::project2} ::proj_1_and_2_doc
```

Because you are copying the files to a special documentation directory, the .DOC type extensions are redundant, and you may remove the extensions from the restored files with this command:

```
co 8{@.doc::project1 @.doc::project2} @.::proj_1_and_2_doc
```

Example 6: Restore the single file X.DOC from the PROJECT1 directory to the file X_DOCUMENT in the working directory.

```
co 8{x.doc::project1} x_document
```

Example 7: Restore all documentation files to the global directory DOCUMENTATION.

```
co 8{@.doc} ::documentation
```

Example 8: Restore all DOCUMENTATION files to FMGR cartridge DO, removing the .DOC type extensions.

```
co 8{@.doc} @.::do
```

Names longer than six characters are truncated.

- Special Purpose Backup Examples

Example 1: Back up all files in global directory PROJECT1, but label the files on tape as if the files came from directory PROJECT2.

```
co ::project1 8{::project2}
```

Example 2: Back up the file X.DOC in the directory PROJECT1, renaming the file INFO_ABOUT_X.

```
co x.doc::project1 8{info_about_x}
```

Example 3: Back up all FORTRAN source files in the working directory, removing the type extensions.

```
co @.ftn 8{@.}
```

Copy Examples with Subdirectories

The examples in this section illustrate how to use subdirectories to take advantage of the hierarchical features of the file system. Many of the points discussed in the previous section “Copy Examples without Subdirectories” also apply when you use subdirectories. Features that are common to both discussions (such as examples of how to use the GR command) are handled in more detail in the previous section.

These examples parallel the discussion in the earlier section “Copy Command Source and Destination Parameters”.

The tape LU in all these examples is LU 8.

- Normal Backup Examples

Example 1: Back up all files in the global directory MYFILES, including all files in all subdirectories of MYFILES, and all subdirectories of those subdirectories, with the Verify option selected.

```
co /myfiles/@.@ 8 V
```

Alternatively, you may select the following equivalent notation; however, the “::” notation may become confusing when subdirectories are specified, as shown below.

```
co @.@::myfiles 8 V
```

When the name and type extension are both @ and no qualifiers are used, you can drop the @.@ to save typing. The commands shown above can thus be abbreviated as follows:

```
co /myfiles/ 8 V
```

or

```
co ::myfiles 8 V
```

TF

Example 2: Back up files in directory MYFILES, subdirectory PROJECT1.

```
co /myfiles/project1/ 8
```

This command backs up all the files in the PROJECT1 subdirectory of global directory MYFILES, as well as all files found in any subdirectories within the PROJECT1 subdirectory.

Example 3: Back up files in subdirectory SRC, within subdirectory PROJECT1, within global directory MYFILES.

```
co /myfiles/project1/src/ 8
```

Alternatively, you could use the “:” notation in this command, as follows:

```
co project1/src/::myfiles 8
```

In this notation, the hierarchical structure is less clear than it was in the first command. It is much easier to see that /MYFILES/PROJECT1/SRC/ refers to a directory MYFILES, which contains the subdirectory PROJECT1, which in turn contains a subdirectory SRC. Therefore, the remaining examples in this section use the hierarchical notation shown in the first command, rather than the “:” notation.

Example 4: Back up all FORTRAN source files in subdirectory SRC.

```
co /myfiles/project1/src/@.ftn 8
```

If you set the working directory to /MYFILES/PROJECT1, you may use the following command to back up the same files:

```
co src/@.ftn 8
```

If you set the working directory to /MYFILES/PROJECT1/SRC, you may use the following command to back up the same files:

```
co @.ftn 8
```

However, when you select files for backup relative to the working directory, the files are identified on the tape with file descriptors such as MODULE1.FTN (or SRC/MODULES1.FTN, for the previous example), which do not precisely specify the directory from which the file was backed up.

When you restore files from such tapes without specifying a destination (as in the restore command “co 8”), the file descriptors on the tape are interpreted relative to whatever working directory is in effect at the time of the restore.

This is convenient in some cases, such as when you transfer files that do not need to reside in a particular directory. However, for ordinary backup applications, we recommend that you specify files with full file descriptors, rather than use the working directory. This ensures that the files are restored to the proper directories.

Example 5: Back up all FORTRAN source files found in subdirectory SRC, or any subdirectory of that subdirectory, with the S mask qualifier specified.

```
co /myfiles/project1/src/@.ftn.s 8
```

If you restore files from this tape to a specified destination directory, all the FORTRAN source files on the tape are copied directly into the directory FILES_FROM_8 without any intervening subdirectories, whether they originally came from SRC or from a subdirectory of SRC. This is due to the S mask qualifier. The command entry is as follows:

```
co 8 /files_from_8/
```

If you restore the files without specifying a destination, all the files are restored to their original directory or subdirectory, as in the following command:

```
co 8
```

Example 6: Back up all files in directory GLOBAL1, subdirectory SUB1, that were not accessed since May 1983.

```
co /global1/sub1/@.a-8305 8
```

Here the A mask qualifier selects files with particular access times (in this case, year 83, month 05).

You can use this command to back up old files before purging them to save disk space.

Example 7: Back up all files in the directory MYFILES, subdirectory PROJECT1, or in any subdirectories contained within the subdirectory, that were not accessed since Jan 4, 1983.

```
co /myfiles/project1/@.sna-830104 8
```

Here the A mask qualifier is used with the S and N qualifiers. A selects files with particular access times (in this case, year 83, month 01, day 04). The N qualifier (No directories) is needed to keep subdirectory files from being selected, which would automatically cause files within them to be selected regardless of their access time.

Example 8: Back up all files in the working directory and in all subdirectories contained within the working directory.

```
co @.@ 8
```

Example 9: Back up FILE1, contained in the working directory.

```
co file1 8
```

Example 10: Back up all files in SUB1, a subdirectory of the working directory, and all files in subdirectories of that subdirectory.

```
co sub1/ 8
```

TF

Example 11: Back up all Macro source files in the working directory.

```
co @.mac 8
```

Example 12: Back up two specific files.

```
co {/global1/file1 /global2/sub1/suba/filex} 8
```

Example 13: Back up all FORTRAN and Macro source files in directory MYFILES, subdirectory PROJECT1, subdirectory SRC.

```
co {/myfiles/project1/src/@.ftn /myfiles/project1/src/@.mac} 8
```

If the working directory is set to /MYFILES/PROJECT1/SRC, you may use the following command to copy the same files to tape:

```
co {@.ftn @.mac} 8
```

However, in this case, the files are identified on the tape with shortened file descriptors that do not specify the directory of origin, rather than full file descriptors such as /MYFILES/PROJECT1/SRC/MODULE1.FTN.

To shorten the command syntax without sacrificing full file descriptors, use the DE command to specify a default directory for the source parameter of the copy command, as follows:

```
de /myfiles/project1/src/  
co {@.ftn @.mac} 8
```

DE specifies defaults for source, destination, and option field values that you do not specify in subsequent commands. In the command above, a default is used for the source directory. The final / after SRC in the DE command indicates that only the files contained in directory SRC are to be backed up, not the directory and its files.

The DE command lets you specify CO commands in abbreviated form. All abbreviations are “expanded” before the command is processed, and do not cause the file names to be abbreviated on the tape. For example, the abbreviated CO command above is expanded into the following command:

```
co {/myfiles/project1/src/@.ftn /myfiles/project1/src/@.mac} 8
```

Example 14: Back up all files in the sources and relocatables subdirectories within the PROJECT1 subdirectory.

```
co {/myfiles/project1/src/ /myfiles/project1/relocatables/} 8
```

- Normal Restore Example

Example 1: Restore all files from tape LU 8.

```
co 8
```

If you do not specify a destination parameter, all the files are restored to the same directories from which they were backed up, with the same names and properties.

However, if files are backed up relative to the working directory, the files are restored to the current working directory in effect when they are restored, which may not be the same as the working directory in effect when they were backed up.

- Selective Restore Examples

Example 1: Restore all files from tape LU 8 that were backed up from subdirectory /MYFILES/PROJECT1/SRC.

```
co 8{/myfiles/project1/src/}
```

This command does not work if the files were backed up relative to the working directory, because in that case, the files are not identified on the tape with the full file descriptor (starting with /MYFILES). Rather, they are identified with a shorter file descriptor. The file mask used in selecting files from tape must reflect this.

You can use the DL command to determine how the files are identified on the tape, so that you know how to select files being restored. (See the section “Relation of the DL command to the CO command” later in this chapter for details.)

Example 2: Restore all files from tape LU 8, selecting only the FORTRAN source files (type extension .ftn).

```
co 8{@.ftn.e}
```

The E qualifier (look Everywhere) must be specified to select files with type extension .FTN at any directory level. If you specify a particular directory, the E qualifier is not necessary, as shown in the following command:

```
co 8{/myfiles/project1/src/@.ftn}
```

Example 3: Restore two specific files.

```
co 8{/global1/file1.ftn /global2/suba/filex.ftn}
```

Example 4: Restore all FORTRAN and MACRO source files from the subdirectory /MYFILES/PROJECT1/SRC.

```
co 8{/myfiles/project1/src/@.ftn /myfiles/project1/src/@.mac}
```

You may use DE to simplify the command, as follows:

```
de /myfiles/project1/src
co 8{@.ftn @.mac}
```

Example 5: Restore all FORTRAN, Pascal, and MACRO source files, regardless of the directory they came from, using the E mask qualifier.

```
co 8{@.ftn.e @.pas.e @.mac.e},,v
```

TF

- Restoring to Different Directories or Files Examples

This type of restore can be selective (specifying a source mask) or non-selective (no source mask). Examples 1-5 illustrate selective restore.

Example 1: Restore all files backed up from global directory /MYFILES to the subdirectory /OLDPROJECTS/CATEGORY1.

```
co 8{/myfiles/} /oldprojects/category1/
```

The CO command above retains the hierarchical structure of the file system. When you copy the contents of global directory MYFILES, all levels of subdirectories are copied intact into the subdirectory CATEGORY1. The subdirectories are automatically created when you restore the files. For example, if you backed up /MYFILES/SUBFILES/FILE1, it is restored with the following path name:

```
/oldprojects/category1/subfiles/file1
```

Because MYFILES itself was not copied (the trailing slash in the source parameter /MYFILES/ says copy only the contents of the directory), that directory is no longer part of the FILE1 path name. Refer to the description of the CO command in the *RTE-A User's Manual* for further information.

Example 2: Restore all documentation files (type extension .DOC) that were backed up from subdirectory PROJECT1 of global directory MYFILES, putting all the files in the global directory DOCUMENTATION.

```
co 8{/myfiles/project1/@.doc} /documentation/
```

Example 3: Restore all documentation files to the DOCUMENTATION directory, regardless of what directory or subdirectory they were backed up from, using the E mask qualifier.

```
co 8{@.doc.e} /documentation/
```

This command causes all the selected files to be copied into the specified directory without any intervening subdirectories, regardless of their original directory or subdirectory. In other words, the E qualifier causes the hierarchical structure to be compressed into one level.

However, when you select a directory file, the hierarchy below the selected directory is preserved. (In the example above, specifying the type extension .DOC excludes directory files.)

Levels are similarly compressed when you use the S qualifier, as follows:

```
co 8{/myfiles/@.doc.s} documentation/
```

This command selects all documentation files in global directory MYFILES (on the tape) or any subdirectory contained within MYFILES. All the selected files are copied into the DOCUMENTATION global directory with no intervening subdirectories, regardless of the directory or subdirectory level from which the file originally came.

Example 4: Restore all the same files and remove the .DOC type extension from the resulting files.

```
co 8{/myfiles/@.doc.s} /documentation/@.
```

Example 5: Restore a particular file, /MYFILES/PROGRAM_X.DOC, from tape into a particular destination file on disk, /DOCUMENTATION/INFO_ABOUT_X.

```
co 8{/myfiles/program_x.doc} /documentation/info_about_x
```

In a non-selective restore, the entire contents of the source tape are restored to the named destination. The path names of the restored files are determined by the way in which you specified the initial copy source parameter. Examples 6–8 illustrate non-selective restore.

Example 6: Back up the contents of SUB2 of SUB1 of MAS.

```
co /mas/sub1/sub2/ 8
```

The directory for this backup tape lists the path followed to access the files, then itemizes the files backed up. For example, assume the tape contains file FILEZ and a subdirectory of SUB2 called SUB3, which contains a file FILEX. You can restore this tape to disk with the following command:

```
co 8 /dira/suba/
```

In this case, FILEZ is restored as /DIRA/SUBA/FILEZ, and FILEX is restored as /DIRA/SUBA/SUB3/FILEX.

Example 7: Back up the MAS directory, all its subdirectories, and all the files therein.

```
co /mas.dir 8
```

You may restore this tape to disk with the command:

```
co 8 /dira/suba/
```

In this case, FILEZ is restored as /DIRA/SUBA/MAS/SUB1/SUB2/FILEZ and FILEX is restored as /DIRA/SUBA/MAS/SUB1/SUB2/SUB3/FILEX.

Note, however, that if you restore the tape with the following command, you delete the directory /MAS from the path name for the restored files:

```
co 8 /dira/suba
```

Example 8: Back up all the files with a .DOC extension in the MAS directory, using the S qualifier.

```
co /mas/@.doc.s 8
```

When you use the S or E qualifiers in selecting the files for backup, TF discards the directories and subdirectories it searched to determine the destination when you restore them with the command:

```
co 8 /dira/suba/
```

the restored files are put in the specified subdirectory with no intervening subdirectories.

You can use the DL command to see the directory, to find out how the files are restored. Refer to the discussion of the DL command in this chapter for more information.

TF

- Special Purpose Backup Examples

Example 1: Back up all files in subdirectory /MYFILES/PROJECT1/SRC, but label the files on tape as if they had come from a global directory called PROJECTX.

```
co /myfiles/project1/src/ 8{/projectx/}
```

Example 2: Back up all the files in the DOCUMENTATION directory, but label the files on tape as if they came from a directory called PROJECT1. Give all the files on tape the type extension .DOC, regardless of what type extension they had on disk.

```
co /documentation/ 8{/project1/@.doc}
```

Example 3: Back up all the system files that have type extension .DOC, label them all as if they came from the global directory DOCUMENTATION, and remove all the .DOC type extensions from the files on tape.

```
co @.doc.e 8{/documentation/@.}
```

Copy Examples Using DS

Although TF cannot access a remote tape unit, you can run TF from a system with a local tape unit and then access the backed up or restored disk files over DS. For remote backups and restores of several files, you must specify a superuser logon name in the command. (Being logged on as a superuser locally does not affect remote access.)

Unlike other file attributes, the DS location for a file being restored does not default to the same location from which the file was backed up. If the destination DS location is not specified, a local destination is assumed.

Be sure that the system times are correctly set at both the local and remote systems. Refer to the section “Maintaining the System Time” later in this chapter for details on the system time stamps. Refer also to the section “Saving and Restoring Properties of Files” later in this chapter for additional information on using TF with remote files.

When files are backed up and restored using DS, delimiters are used to specify a remote node and logon. Note that if “]” or “>” occurs at the beginning of a file descriptor or immediately following a slash, it is treated as the first character of a file name, not as a DS delimiter.

When you use a DS delimiter to specify a remote destination, you must precede the DS delimiter by at least one other character.

Example 1: Restore a tape to the same directories from which the files were copied at DS node >mynode.

```
co 8 0>mynode
```

If the delimiter is the first character, precede it with a “0.” TF interprets the zero as a blank and correctly interprets the next character as the DS delimiter. The following command is incorrect:

```
co 8 >mynode
```

In this case, TF interprets the > delimiter as the first character of a file name, and it restores the tape to local file >mynode.

Example 2: Restore tape containing a single file to filez at DS node >mynode

```
co 8 filez>mynode
```

In this case, the file name filez precedes the DS delimiter >, and the backup is successful.

When you specify a remote file descriptor in the source or destination, you must not specify the delimiter as the first character following the “/” in the descriptor. Use an “at” sign (@) after the slash to identify the > or [as a DS delimiter. The next example illustrates the use of the @ sign.

Example 3: Back up files in /DIRA at DS location [mylogon]>mynode to LU 8:

```
co /dira/@[mylogon]>mynode 8
```

In this command, the delimiter is preceded with @. This tells TF that the next character is a DS delimiter, rather than the first character of a file name. The following command is incorrect:

```
co /dira/>mynode 8
```

Here TF interprets > as part of the file name and attempts to back up the local /DIRA/>mynode files to tape.

Example 4: Restore tape to /DIRA at remote location [mylogon]>mynode, where [mylogon] is the user logon name at the DS location.

```
co 8 /dira/@[mylogon]>mynode
```

Note that the logon name is used only in multiuser systems. If the DS location is not multiuser, you must not use a logon name. If a password is associated with the logon name (multiuser systems only), you must specify it. For example:

```
co /dira/@[mylogname/mypassword]>mynode 8
```

With the exception of DS delimiter handling, copy operations with TF are identical for both local and remote backups and restores. The following examples summarize the similarities and differences in the TF command syntax. In all cases, the first command in each example is a local TF command, and the second is a remote TF command.

Example 5: Back up all files in global directory MYFILES.

```
co /myfiles/ 8
co /myfiles/@>mynode 8
```

Example 6: Restore all files to global directory THEFILES.

```
co 8 /thefiles/
co 8 /thefiles/@>thenode
```

Example 7: Back up all files in the hierarchical file system (system file backup, generally done only by a superuser).

```
co / 8
co /@[superusername]>thenode 8
```

TF

Example 8: Restore all files to their original directories.

```
co 8  
co 8 0[superusername]>thenode
```

Example 9: Back up all files on file system LU 27.

```
co 27 8  
co 27[superusername]>thenode 8
```

Example 10: Restore all files to file system LU 54 (except for files in directories that already exist on other LUs).

```
co 8 54  
co 8 54[superusername]>thenode
```

In all the examples above, @ is used only in the remote command because it is a requirement when using TF with DS; however, you may specify @ in the local file descriptors if desired. The three commands that follow all have the same effect in a local backup of all files from global directory MYFILES:

```
co /myfiles/ 8  
co /myfiles/@ 8  
co /myfiles/@@ 8
```

For remote node operations, only the following two forms of the command work properly:

```
co /myfiles/@[user]>node 8  
co /myfiles/@@[user]>node 8
```

Default (DE)

Purpose: To set up defaults for the CO command source, destination, and options parameters.

Syntax: DE *source destination options*

source As defined for the CO command, except that a
destination list of masks may not be used in the source parameter.
options

Description:

The DE command is useful when you are entering a number of CO commands in succession, and all or most of them have the same parameters, such as the same source or destination tape LU or disk directory, or the same options. DE lets you specify the value(s) once, for all the related CO commands.

The DE command is especially useful in conjunction with the GR command (see the section on “Group Copy Commands” and the example of DE command use in a subsequent section).

You may use the DE command source parameter to set up default values for the tape LU, file name, type extension, file mask qualifier, security code, directory (specified in the directory field or in front of the file name), type, size, record length, or DS location items of the CO command source parameter. The default value is subsequently used whenever you do not specify another value in the source parameter of the CO command, except for the qualifier field. Mask qualifiers included in the DE command are merged with any qualifiers specified in subsequent CO commands.

Similarly, the DE command destination parameter can be used to set up default values for the tape LU, name, type extension, security code, directory, size, or DS location items of the CO command destination parameter.

In order to default the DS location only in the source or destination parameter, enter a colon (:) before the “[” or “>” character at the beginning of the DS field. This keeps the “[” or “>” from being interpreted as the first character of the file name. Note that in these examples, “[”, “]”, and “>” are characters that are actually typed, rather than part of the syntax notation. For example:

```
de :[sourceuser] :[destuser]
de :>sourcnode :[destuser]
de :[sourceuser]>sourcnode :>destnode
```

Note that the user name and node name are both part of the DS location and are not separated when defaults are used. If the CO command source (or destination) parameter specifies either the user name or the node name, the default value is not used for either name.

When you include options in the DE command, they are merged with any options you specify in subsequent CO commands. For example, a DE/CO command set that has the following form:

```
de , , ,pv
co source destination cp
co source destination
```

specifies options C, P, and V for the first CO command, and P and V for the second.

Each DE command supersedes the preceding default; therefore, you can cancel defaults by entering a DE command with all null parameters.

TF

Directory List (DL)

Purpose: Provides a directory list of files on a TF tape specified by the command.

Syntax: DL <tape> [I][K]

or

DL <tape>{<mask>} [I][K]

or

DL <tape>{<mask1> <mask2> ...} [I][K]

tape The positive or negative tape LU.

mask The specifier of certain files to be listed.

Description:

Because the mask works exactly the same as in the CO command, you can use DL to show what files are to be copied by a particular mask. If you do not specify a mask, all the files on the tape are included in the listing.

The I option causes TF to ignore errors in the tape header and unexpected file marks. Refer to the earlier description of the I option for more details.

You may specify the K option to keep TF from taking the tape unit offline. Be aware that when you use K, the tape unit is unlocked but left online; thus, it is possible for another program to overwrite the tape. Refer to the section in this chapter on “Tape Protection” for more information.

You can use LL to direct the DL listing to another device or disk file.

The following example shows the directory list format:

```

TF:                               dl 8

Tape format:                       TF
Title:                             TF:  co {/programs/ /project1/} 8 v
Date:                              Mon Mar 21, 1983  5:21:59 pm

Created Mon Mar 21, 1983  5:21:59 pm:

File from directory /PROGRAMS/      Update Time or Owner

FOWN.RUN:::6:123:128                Tue Mar 15, 1983 11:38:04 am
FREES.RUN:::6:39:128                Tue Mar 15, 1983 11:38:06 am
FPACK.RUN:::6:158:128               Tue Mar 15, 1983 11:38:10 am
CLOSE.RUN:::6:12:128                Tue Mar 15, 1983 11:38:12 am
CI.RUN:::6:229:128                  Tue Mar 15, 1983 11:38:21 am
FVERI.RUN:::6:152:128               Tue Mar 15, 1983 11:38:08 am
DL.RUN:::6:136:128                  Tue Mar 15, 1983 11:38:03 am

File from directory /PROJECT1/      Update Time or Owner

SRC.DIR:::2::32                      MYUSERNAME
INCLUDES.DIR:::2::32                 MYUSERNAME
SRC/MODULE1:::4:31:39                Fri Mar  4, 1983  3:32:32 pm
SRC/MODULE3:::4:136:39               Thu Mar  3, 1983  1:52:36 pm
SRC/MODULE2:::4:236:39               Fri Mar  3, 1983  3:32:50 pm
SRC/MODULE4:::4:20:38                Fri Mar  4, 1983  3:32:51 pm
SRC/MODULE6:::4:2:34                 Fri Mar  4, 1983  3:32:53 pm
SRC/MODULE5:::4:397:39               Fri Mar  4, 1983  3:33:00 pm
INCLUDES/INCLUDE1:::4:1:31           Wed Mar  2, 1983 10:27:12 am
INCLUDES/INCLUDE2:::4:1:17           Wed Mar  2, 1983 11:14:34 am
INCLUDES/INCLUDE3:::4:4:36           Fri Mar  4, 1983  3:33:16 pm
INCLUDES/INCLUDE4:::4:12:35          Tue Mar  1, 1983  4:15:25 pm

```

Figure 2-6. Example of Directory List Format

The directory listing begins with the tape header. The date in the header is the date and time at which the first file was written to the tape.

The “Created” line represents the date and approximate time at which the following list of files was created on the tape by the copy operation; it does not represent the time at which the original file was created. Whenever files are appended to the tape, a new “Created” line precedes the list of the new files and gives the date and time they were appended.

If the “File from directory” line appears, the directory name combines with listed files to form the full file descriptor (for example, /PROGRAMS/FOWN.RUN:::6:123:128). If this line does not appear, the file as listed is the full file descriptor.

Note that the update time for FMGR files is always reported as 1/1/70 because FMGR files do not have an update time in their directory entry.

TF

Relation of the DL Command to the CO Command

The DL command helps you determine how to select files to copy from a tape. You may select a file by its full file descriptor, which includes the name from the directory heading line. Alternatively, you may select one by the reference file descriptor on the file line that follows the directory heading line (if a directory is specified in the line). Referring to the example directory listing shown in the DL command section above, the following two CO commands select the same file from the tape:

```
co 8{/project1/src/module1}      (Full file descriptor)
co 8{src/module1}                (Reference file descriptor)
```

(If the header line does not specify a directory, both file descriptors are identical.)

If you do not know which files a given file mask selects, try the mask first with DL before you use it in a CO command. For example:

```
dl 8 {src/module1}
```

When you issue this command, the files listed are the ones that will be copied if you use the same mask with the CO command.

You can use the reference file descriptor to determine the directory hierarchy created by a non-selective restore to a specified directory, as in the following command:

```
co 8 /global/sub/
```

Assume the tape is the same as the one shown in the example of the DL listing above. The first file that this command restores is

```
/GLOBAL/SUB/FOWN.RUN    (not /GLOBAL/SUB/PROGRAMS/FOWN.RUN)
```

because the reference file descriptor, rather than the full file descriptor, is used.

Because this command does not copy the programs directory to tape but just the files in it (it does not specify /PROGRAM.DIR, just /PROGRAMS/), no programs directory is created when the files are copied from the tape. If you specify /PROGRAMS.DIR in the backup command, rather than /PROGRAMS/, the reference file descriptor is /PROGRAMS/FOWN.RUN instead of FOWN.RUN, and the file is restored as /GLOBAL/SUB/PROGRAMS/FOWN.RUN.

The effect of using this method is that when you back up files and later restore them to a different destination, as in the following two commands:

```
co /programs/ 8
co 8 /global/sub/
```

the result is the same as if you copied the files directly from the source in the first command to the destination in the second command, as in the following CI utility command:

```
CI> co /programs/ /global/sub/
```

The directory header line always shows the name of the global directory (or FMGR cartridge), and the file lines include the name and type extension. Files from such tapes can be selected by any combination of properties, for example:

```
co 8{fown.run}
co 8{fown.run::programs}
co 8{:::programs}
co 8{@.ftn::project2}
```

Exit (EX)

Purpose: Terminates TF.

Syntax: EX [K]

K This option keeps the most recently used tape LU online. Note that this makes it possible for another program to overwrite the tape, because the tape unit is unlocked, but left online. Refer to the section on “Tape Protection and the K Option” for more details.

Group Copy Commands (GR, EG, and AG)

Purpose: To combine multiple CO commands into a single operation.

Syntax: group
 co *parameters*
 co *parameters*
 .
 .
 .
 co *parameters*
 eg

Description:

Some copy applications require combining multiple CO commands into a single operation. It is useful to group commands, for example, when there are too many source masks to fit on one command line, or when different destinations or options are needed for different files.

When TF encounters a GR command, it keeps track of subsequent CO commands, but does not execute them immediately. Instead, TF executes them as a single operation when it encounters the EG (End Group) command. You can abort the GR operation by using the AG (Abort Group) command, which causes the group to be terminated before any of the CO commands are executed.

If the grouped CO commands have a common source or destination, you may use the DE (Default) command to advantage. For example:

```
de 8                (Common source tape LU)
gr                 (Start group)
co ::x ::a        (Restore files from directory X to directory A)
co ::y ::b        (Restore files from directory Y to directory B)
co ::z ::c        (Restore files from directory Z to directory C)
eg.               (End group)
de,,8             (Common destination tape 8)
gr                 (Start group)
co abcd wxyz      (Copy file ABCD to 8 as WXYZ)
co efghij qrstu   (Copy file EFGHIJ to 8 as QRSTU)
eg                 (End group)
```

Refer also to the earlier section on the DE command.

If you group many CO commands, you may need to increase the size of EMA for TF, otherwise the following message may be displayed:

```
EMA size of TF not large enough for a group of this size.
```

See the section on “Installing TF” later in this chapter for more information.

TF

Help (HE)

Purpose: Writes a summary of TF commands to the terminal.

Syntax: HE or ?

List Header File (LH)

Purpose: Lists the tape header file to the list device.

Syntax: LH *tape* [K]

tape The positive or negative tape LU.

K The option that keeps TF from taking the tape offline. Note that this option makes it possible for another program to overwrite the tape, because the tape unit is unlocked but left online. See the section on “Tape Protection” for more information.

Description:

The following is an example of the LH command listing:

```
Tape Format:    TF
Title:         TF: co /project1.dir 8
Date:         Mon Mar 14, 1983  1:28:05 pm
Capacity:     16352 kilobytes
Used:        725  kilobytes
```

“Capacity” and “used” lines appear only for CTD tapes and indicate the total data capacity of the CTD cartridge and the amount of data that the cartridge currently contains. (A CTD block is one kilobyte.)

For multi-tape backups, tapes that follow the first one also show the tape number in the LH listing after the title and date (for example, “Tape number: 2”).

List Device (LL)

Purpose: To specify a device or file other than the terminal to receive list information.

Syntax: LL *file*

file The file descriptor of the device or disk file to receive the list information. If a disk file is specified, it is created. An existing disk file is not overwritten.

Description:

Initially, the list device for TF is the terminal.

Tape directory listings, header files, comment files, “?” command information, and command echoes are listed with the LL command. Information and error messages are always displayed on the terminal and are not affected by the LL command.

Title (TI)

Purpose: Sets the title to be used on tapes created by subsequent CO commands.

Syntax: TI *title*

title The statement of up to 80 characters (extra characters are truncated and a warning is given) that identifies a particular tape.

Description:

The TI command must precede the CO (or GR) command that defines the files to be copied.

You can use the LH command to read the title of a tape. In addition, the title appears at the beginning of a DL command listing and is displayed when you copy files from a tape. The title appears first on a TF tape, and for magnetic tape, can be listed directly from the CI LI (List File) command. The following example shows how to specify a title when you back up files and how to list the title using LH or LI.

```

CI> tf
TF: ti FORTRAN source files for project XYZ.
Title used will be: FORTRAN source files for project XYZ.
TF: co /xyz/@.ftn 8 v
:
:
TF: lh 8
Tape format: TF
Title: FORTRAN source files for project XYZ.
Date: Thu Mar 17, 1983 10:35:41 am
TF: ex k
CI> li 8,,1
FORTRAN source files for project XYZ.
CI>

```

Note that when LI is used, other information may follow the title on subsequent lines. This information can be ignored.

If you enter a CO command but do not specify a title, TF supplies a default title, which consists of the CO command (exactly as typed), preceded by "TF:". In the example above, if you omit the TI command, the tape title is:

```

Tape format:      TF
Title:           TF: co /xyz/@.ftn 8 v
Date:           Thu Mar 17, 1983 10:35:41 am

```

"TF:" is included in the title to make the tape easier to identify when it is listed directly, such as when you use the CI LI command.

TI has no effect when you append files to a tape (that is, when you select the Append option to use with the CO command).

TF

Transfer Command (TR)

Purpose: Lets you transfer control to and return from a command file.

Syntax: TR *file* | TR

file The command file or input device to which control is being transferred.

TR Return to the next higher level command file when TR files are nested. (When encountered in the first level command file, this serves the same function as the TF EXIT command.) A return is done automatically when the end of a command file is reached, so this form of the TR command is not usually needed.

Description:

Command files can be nested to four levels, with control passed from file level to file level using the TR command.

Tape Protection and the K Option

TF keeps the tape LU locked when you issue commands that access tape to protect the tape against access by other programs. If the tape is write-protected, TF unlocks the LU at the end of the command because no further protection is necessary.

If the tape is not write-protected, TF keeps the tape LU locked until the utility exits or another tape LU is accessed. When the LU is unlocked, the unit is taken offline (in the case of the CTD, the cartridge is unloaded by TF) to protect the data.

At times, you may need to access the same tape on more than one successive TF command (such as LH followed by DL). In those cases, enter all the commands as a single execution of TF, as follows:

```
CI> tf
TF: lh 8
TF: dl 8
TF: ex
```

In the example above, the tape is not taken offline until the EX command is encountered.

Write-protecting the tape also keeps it from being taken offline.

Another way to keep the tape online is to use the K option. However, K leaves write-enabled tapes unprotected from the possibility of being overwritten by other programs. The K option can be specified with any tape command (CO, DL, or LH) or in the EX command. The option need only be specified once each time TF is run, because it remains in effect until TF terminates.

Be particularly careful to protect a tape being appended because the data already on the tape must be preserved, yet the tape cannot be write-protected. You should load the tape after you enter the CO command rather than before. (For CTD, all that matters is that the tape does not complete the loading process until after the command is entered). When you enter CO, TF displays the message:

```
Tape not ready.
Type GO when tape is ready, or type BR to terminate.
```

Load the tape, then type GO. This ensures that the tape LU is locked before the tape is ready for access, which protects the previous contents of the tape from being overwritten.

Time Stamps

The file system maintains three time stamps for each file on disk: create time, update time, and access time. Similarly, TF maintains time stamps for files on tape. The meanings of time stamps on tape are analogous to their meanings on disk.

Note that update times are preserved when files are copied (whereas create times and access times are not). This is done so that the update time can be used as an indication of the revision date of the file. It is important that the revision date is not lost when a file is copied, backed up, or restored.

Create Time

This is the approximate time the file was created on the tape. The same time is used for all files copied to tape in the same command (or group of commands) so that separate segments of incremental backups can be easily identified. This time stamp is not present on UNIX TAR tapes. (See “Missing Time Stamps” below.)

Update Time

Because the file on tape is never directly updated, but is always just a copy of a disk file, the update time for the tape file is always the same as the update time of the disk file being copied.

Access Time

No access time is maintained for tape files.

Missing Time Stamps

As indicated above, files on tape do not have access time stamps, and files on UNIX TAR tapes also do not have create time stamps.

For the purposes of selecting files from a tape and doing a DL of a tape, it appears as if these missing time stamps are set to 1/1/70. This is the time that is used by convention whenever the true time is unknown or undefined, such as in FMGR files.

Maintaining the System Time

Confusing and perhaps serious consequences may occur if you do not maintain the system time (or set it backward during a restore). Some of the possibilities are as follows:

- The Update option may produce undesirable results if the system time is not maintained correctly.
- Even if you did not specify the Duplicate or Update option, duplicate files may be purged without warning on a restore.
- Restoring incremental backups may not work correctly if the system time is set backward during a restore. TF may restore out-of-date versions of some files rather than the latest version.

When TF is used to access remote files (that is, over DS), the system time must be correct at both the local and the remote system.

TF

Incremental Backup

Incremental backup is a procedure that involves periodically backing up all the files and doing frequent backups only of files that were changed since the previous backup. (This procedure is not applicable for FMGR files.) In this discussion, a backup of all files is called a “full backup,” and a backup of just those files that were changed is called a “delta backup.”

An incremental backup consists of one full backup followed by a series of delta backups. Restoring from an incremental backup, therefore, consists of restoring the most recent full backup followed by all of the delta backups done since then.

In the standard incremental backup procedure, a full backup is done at the beginning of a tape, and then subsequent delta backups are appended to the same tape. Each new incremental backup is done on another tape. (An alternative to the standard incremental backup procedure is discussed below.)

Incremental backup has the following advantages over ordinary backup:

- The average backup time is faster because most of the backups are delta backups, which generally take less time than full backups. As a result, system availability is improved. Because the incremental backup process is faster, other access to the files is restricted for a shorter period of time. (See the section in this chapter on “File Access During Backup and Restore Operations.”)
- Less tape is used on the average.
- Fewer tapes are used because several backups are appended to the same tape.
- Multiple versions of files can be accessed more conveniently on tape (good for archiving applications).

The disadvantages are that incremental backups take longer to restore, and the procedures involved are more difficult to understand.

Backup Bit, B Qualifier, and C Option

Incremental backup is accomplished by making use of the backup bit. The file system provides a backup bit for each file (excluding FMGR files). Because the backup bit is automatically set by the file system whenever a file is modified, it can be used to keep track of which files were changed since the last time they were backed up.

The backup bit is cleared when a file is backed up with the Clear Backup Bit (C) option specified. The Clear option also implies the Verify option. The backup bit for each file is cleared after the file is verified. This ensures that a file’s backup bit is not cleared until the file is successfully backed up.

If the B qualifier is specified on a backup (as in the file mask `/PROJECT1/@.@.B`), the backup bit is checked to determine which files are to be backed up. Using this procedure, it is possible to back up only those files that were changed since the previous backup. This does not apply to FMGR files, as specifying the B qualifier for those files has no effect and all the FMGR files are copied.

To ensure that incremental backup works correctly under all circumstances, a fairly complex scheme is used in handling the backup bit for a file and interpreting the B mask qualifier. The file system sets a file's backup bit when any of the following conditions occur:

- The file is created.
- The file's contents are changed.
- The file's name, type extension, protection bits, or owner (directory files only) is changed.
- The file is moved to a different directory (by the CI MO command or by a call to FmpRename).

A file's backup bit is cleared when it is copied to tape by TF with the Clear option specified.

To determine whether a given file is selected by a file mask containing the B qualifier, TF first considers all other fields of the file mask to learn whether the file matches the mask (the B qualifier is ignored). If a directory file matches the mask, all the files in that directory automatically match (due to the implicit D qualifier used in copy commands).

If a file matches the mask this way, either directly or because its parent directory matches, then TF considers the backup bit to determine whether the file will actually be selected. This depends not only upon the backup bit for the file itself, but also upon the backup bits for the directories above the file in the hierarchy up to the level where the mask search started. (For the mask /A/B/@.@.b and the file /A/B/C/D/E, the backup bits in C, D, and E are checked.)

The net result is that a file is selected if the following two conditions are satisfied:

- The file or any of the directories above it match the mask, disregarding the B qualifier.
- The file or any of the directories above it (up to the level where the search began) has its backup bit set.

Standard Incremental Backup Procedure

In terms of options and qualifiers, the standard incremental backup procedure consists of periodically doing a full backup (no B qualifier) at the beginning of a tape (no Append option), and then more frequently (perhaps daily) appending a delta backup (Append option and B qualifier specified). In both cases, the Clear option is specified.

For example, you may perform an incremental backup of the directory /PROJECT1/DOCUMENTATION.DIR (and all files in the directory and all subdirectories) to tape LU 8 as follows:

1. On the first day of the week start with another tape and use this command to do a full backup:

```
co /project1/documentation.dir 8 c
```

2. On each following day of the week, mount the same tape and use the following command to append a delta backup to the tape (note the addition of the B qualifier and the Append option):

```
co /project1/documentation.dir.b 8 ac
```

TF

It is important to use the same source file mask for every segment of an incremental backup (except for the presence or absence of the B qualifier). Otherwise, you may encounter confusing results when you restore the files.

Selecting files relative to the working directory is also discouraged because the working directory may be changed. To reduce typing errors, which could have serious consequences, standard backup commands should be kept in TF transfer files.

Before you decide how often to do a delta backup and how often to start a new incremental backup, there are several factors to consider, such as the frequency of file modifications, the size of the tape reel, and so on.

Note that for CTD, the amount of tape used and the total amount on the cartridge is displayed at the end of each backup, and you can use the LH command to display the amount at any time. For magnetic tape, you can determine the tape usage only by observing the tape during backup.

Multiple Copies of the Same Backup

You can repeat the same CO command to make multiple copies of the same backup (on different tapes). However, for incremental backups, you must omit the Clear option except for the final CO command. When you do this, it is important to make sure that none of the files involved in the backup get updated in between two of the copies. Otherwise, the multiple copies will contain different files, and only the last one (which had the Clear option) will be valid.

Reusing the Backup Bit

If you maintain two independent sets of backup tapes that contain files in common, you cannot use incremental backup for both sets of tapes because there is only one backup bit per file. For example, if everyone backs up his or her own files and, in addition, the system manager backs up everyone's files, incremental backup should NOT be used. Rather than merely providing redundant backups of the same files, in this case incremental backup would result in incomplete backups of all the files.

Restoring Incremental Backups

To restore the entire contents of a backup tape, use the following command:

```
co tape
```

This restores all files on the tape as they were at the time of the backup. If the tape is an incremental backup, this restores the files as they were at the time of the last backup; however, some files that were purged, renamed, or moved elsewhere since the first segment on the tape was made may also be restored.

For example, files that were renamed may be restored both with the old and the new name. Although some old files that no longer exist may be restored, it is always guaranteed that all the files that did exist at the time of the last backup segment are restored correctly.

Incremental backup tapes may contain multiple versions of a given file. When you restore from an incremental backup tape, the latest version of each file is restored. The latest version is the last

copy of the file on the tape, not necessarily the version with the latest update time. (This is because it is possible to discard a new version and revive an old version of a file.) However, in the process of restoring the tape, each version of the file is restored and is replaced by the next version when it is encountered. Therefore, once the entire tape has been restored, only the last version of each file remains.

When you restore incremental backup tapes, it is normal to see the message “Replacing <name> (created during this copy command)” as each version of a file is replaced by the next version. If you specified the Verify option, it is also normal to see the message “Update times don’t match. Not Verifying <file>” as each older version of a file is encountered during the verify process.

The replacement of each version by the next is automatic, and does not require you to specify the Duplicate option. In TF, Duplicate determines whether a file from the tape replaces a duplicate disk file that already exists when the restore operation starts. If a duplicate file does exist, TF assumes that it was created by TF earlier in the restore, and thus is an earlier version of the same file; therefore, TF replaces it without requiring any options.

Although this assumption may not always be true, the results are always correct unless another program creates a file that duplicates a file TF is in the process of restoring. If this happened, TF would inadvertently purge and replace a file created by the other program, even if Duplicate was not specified. To avoid this problem, make sure no other programs access directories that are in use during a TF restore operation.

If a restore operation from an incremental backup tape does not reach completion for any reason, even the files that appear to have been restored correctly may not be up to date, because the latest version of each file is the one that occurs last on the tape.

Restoring Older Versions from Incremental Backup Tapes

To restore files to the state they were in (not as of the most recent backup segment, but as of some previous segment), select files with create times that are less than or equal to the time of that particular segment. For example:

```
co 8{@.@.c-830704}
```

This is the most reliable way to restore an earlier state of the file environment. If you restore the file based on the update time, you may not achieve the results you expect, because update times are not meaningful for directories.

The DL command can be helpful in deciding what to do in such cases. The DL listing groups files according to backup segments. Each segment has a heading that shows the create time for the files in that segment (the time the files were created on the tape).

Setting Backup Bits on Restore

When files are restored, their backup bits are automatically set. This indicates that the newly restored files were not backed up. (Even though the files were just restored from tape, that is not the same as being backed up, because the tape containing the files may not be a backup tape, and may not have been saved as a backup.)

TF

Alternatives to Standard Incremental Backup

An alternative way of backing up files is to start a new tape with a delta backup rather than a full backup. The result is an incremental backup that is not confined to a single tape (or multi-tape set). There are a number of variations of this method. One possibility is to put only one backup segment on a tape. This may be desirable if your delta backups tend to be large.

Another possibility is to do the full backup on one tape and all the delta backups on another tape. This ensures that the full backup tape cannot be damaged during the delta backup. It also becomes practical to keep the delta backup tape on a drive over several backups and devise a method for automatic delta backups at frequent intervals. Although there are risks in keeping a tape on the drive, only the delta backups are at risk.

Because the delta backup tapes are not a multi-tape set, each must be restored with a separate CO command. You must use the Duplicate option with every CO command after the first command, because TF automatically replaces duplicate files only if they are encountered during a single CO operation. (Be sure to restore the tapes in the correct sequence so that duplicate files are replaced in the correct order.)

Replacing Duplicate Files

TF uses certain rules to determine whether to replace a duplicate file that it encounters during a restore operation. A duplicate file is one that already exists with the same name and directory as the file being restored from tape. In making the determination, TF considers whether the Duplicate and Update options were specified and whether the duplicate file was created since the restore operation began.

This applies only to ordinary files, as opposed to directory files. A duplicate directory file is not treated as an error. When a duplicate directory is encountered, the existing directory is left there and is not replaced. Files in the directory may be restored to the existing directory, so there is no need to replace the directory.

TF puts a duplicate file into one of two categories, based on how the create time relates to the time the restore operation started. (This does not apply to FMGR files, which are discussed separately.)

- New duplicate file: It was created since the restore operation began.

TF assumes the file was created earlier in the restore process (by TF itself) and replaces it with the new file unconditionally. This is done so that incremental backups can be restored properly without requiring any special options. The reasons for this and possible side effects are discussed in the earlier section on “Incremental Backup.”

When a new duplicate file is replaced, TF logs the message:

```
Replacing <name>          (created during this copy command).
```

The original file is purged before the tape file is restored. If an error occurs when the file is restored from tape, the original file is still purged but is not replaced by anything. This is done so that errors in restoring incremental backup tapes do not cause the wrong version of a file to be restored. (It is better to restore no file at all than to restore the wrong version.)

- Old duplicate file: It was created before the restore operation began.

Such a file could not have been created by the restore process and is therefore dealt with in the normal way, based upon the D and U options.

If the Duplicate option is specified, the duplicate file is replaced by the file from the tape.

If the Update option is specified, the duplicate file is replaced by the file from the tape only if the update time of the file from the tape is NEWER than the update time of the duplicate file on disk. The error “File already exists” is reported for all other files. This is useful for updating software and similar applications.

When an old duplicate file is replaced, TF logs the message:

```
Replacing <name>
```

When the attempt is made to replace a file because of the Duplicate or Update option, the duplicate disk file is not replaced until after the file from tape is restored successfully to a scratch file, which is then renamed. If an error occurs when the file is restored from tape, the file originally present on the disk remains unchanged.

Because FMGR files have no time stamps, they are never replaced unless the Duplicate option is specified. For FMGR files, you cannot replace duplicate files with security codes, unless the security code of the file you are restoring matches the security code of the existing file. You may use the destination parameter to specify a matching security code, to allow replacement to occur.

Automatic Creation of Directories on Restore

If you try to copy a file into a directory that does not exist, TF automatically creates all the directories required for the files being copied. This can happen when the CO command source parameter specifies the files in a directory, but not the directory itself (for example, you specified “/example/” instead of “/example.dir”), or when the CO command destination parameter specifies one or more new parent directories for the files being copied, as shown in the following command:

```
copy 8{/myfiles/@.ftn} /myfiles/sourcefiles/fortran/@.ftn
```

When a directory file is copied explicitly, its various properties (protection, owner, and disk LU) are sometimes copied as well (see the section “Saving and Restoring File Properties,” below). However, when a directory file is created automatically, there is no file from which to copy the properties, so default values are used.

TF

Saving and Restoring File Properties

TF does more than just save and restore file descriptors and file data. It also preserves the various properties of each file it copies. The following properties are saved for each file copied to tape (excluding directory files) and restored when the file is copied back to disk:

Security code (FMGR files only)

Numeric file type

Record length

Protection information

Update time (See the section on “Time Stamps.”)

For directory files, the following additional information is saved:

Owner name

Disk LU

When remote directories are copied to tape, the true owner is not saved. Instead, all remote directory files are considered to be owned by the user name SYSTEM.

Note that DS location (user name—not to be confused with directory owner and node name) is not saved on tape. To prevent confusion, you should not save files from more than one DS node on the same tape.

Properties of a directory file are only restored when you copy the directory file itself. To ensure that the directory file is backed up, use the command:

```
copy /example.dir <tape>
```

rather than

```
copy /example/ <tape>
```

or

```
copy ::example <tape>
```

which just copy the files in the directory.

If you are copying from tape a directory file that already exists on disk, a “File already exists” error does not occur (as it does for non-directory files). Rather, the existing directory file is used as it is, and other files from tape may be copied into the existing directory.

In this case, the properties of the existing directory are left unchanged. Directory properties are only restored when the directory is created by the restore process, not when a previously existing directory is used.

Normally, the owner of each restored directory file defaults to the user who does the restore. The original owner of each directory is only restored when a superuser does the restore. This facilitates system backup and restore by a superuser. Ownership is not restored for other users, because that causes inconvenience for the routine transfer of files between users (via tape).

Ownership is never restored for remote files. The user specified in brackets in the DS location becomes the owner of all directory files created remotely.

When TF restores a directory file as a new global directory, it creates the global directory on the same LU from which the directory from tape was saved, whenever possible. This is not applicable for subdirectories, which must always go on the same LU as their parent.

If the original LU is unknown (for example, for directories created automatically, or directories restored from UNIX TAR tapes) or unavailable (for example, not mounted or no valid disk LU), the default LU chosen by the file system is used. The file system defaults the LU for creating directories to the same LU as the current working directory, if any. More information about the default LU is provided in the documentation for the “FmpCreateDir” call in the *RTE-A Programmer’s Reference Manual*, part number 92077-90007.

This discussion assumes that the CO command does not specify a destination disk LU. If it does, all global directories (including “automatically” created directories) are created on the specified LU whenever possible. If a CI directory cannot be created on the specified LU, the default LU is used.

System Backup and Restore

This section concerns system-wide file backup and does not cover physical backup of the disk LU that contains the operating system and boot extension files, which is also necessary. You may use the ASAVE/ARSTR utilities to back up and restore the system disk LU. These utilities are discussed in Chapter 3.

You may use the following TF command to back up all files on the system, excluding FMGR files:

```
co / tape
```

The “/” source file mask is an abbreviation for “/@@” which selects all global directories. Because selecting a directory implies selecting all the files in it, this mask selects all files in global directories and their subdirectories. Various options can be added to the above command as required.

You can back up all files on a particular disk LU with the command:

```
co diskLU tape
```

Specifying a disk LU is similar to specifying “/” (in that the root directory is the starting point for the copy), except that only a single disk LU is selected.

You may also perform incremental backup of all non-FMGR files on the system. The section in this chapter on “Incremental Backup” provides a definition of the terms used and complete information about the procedure. The commands to use for full backups and delta backups of all non-FMGR files are presented below for your convenience. Note that you may not select incremental backup for a specific disk LU except by specifying all global directories on that LU.

For full backups, start a new tape and use this command:

```
co /@.@ tape c
```

For delta backups, remount the same tape and use this command, which appends all modified files to the tape:

```
co /@.@.b tape ac
```

TF

Other options can be added to these commands as required.

If you use incremental backup, you must back up FMGR files separately.

If you do not use incremental backup, you may issue the following command to back up all files on the system, including FMGR files:

```
co @.@.e tape
```

The E qualifier means search everywhere, including FMGR cartridges. Be aware of the restrictions that apply when using TF with FMGR files (see the section on “Using TF with FMGR Files” later in this chapter).

Do not use incremental system backup if individual users do incremental backups. Refer to the earlier section in this chapter on “Reusing the Backup Bit”.

The commands shown above for system backup are not restricted to a superuser, but only a superuser will be able to back up read-protected files owned by someone else.

You can restore a system backup tape made by any of the commands discussed above by using the following TF command:

```
co tape
```

You need not specify a destination parameter, because all files are to be restored with the same names and to the same directories from which they were backed up.

Normally, all files are restored to the same LUs from which they were backed up. However, the LUs must be mounted before you do the restore. If you do not need the previous contents of the LUs (which is usually the case when you restore a system), you should clear the LUs prior to the restore (using the CI IN command or the FMGR IN command, as appropriate).

A superuser must do the restore in order to re-establish the original owner of each directory restored.

You may put all the global directories created during a restore on a specific disk LU, by issuing the following command:

```
co tape diskLU
```

If you are entering a directory that already exists on another LU, no error is reported. If a directory cannot be created on the specified LU for any other reason, TF tries to create it on the default LU chosen by the file system. Refer to the earlier section in this chapter, “Saving and Restoring File Properties,” for more details. No error is reported when a different LU is used.

Handling Disk Full Errors

When a disk full error occurs, TF reports the error and suspends operation. The error message contains the name of the file being copied when the disk became full; this lets you determine the LU that is full. You can free space on the disk using MPACK, FPACK, or the FMGR command PK, then restart TF (using the system GO command). Refer to the description of MPACK, FPACK, or PK in the *RTE-A User's Manual* for details of disk packing.

This error may recur during the restore operation, but as long as the CO command eventually completes normally, no files are corrupted or omitted from the copy. However, if the disk full error cannot be corrected, you must enter a system BR TF, then GO TF to break the utility.

File Access During Backup and Restore Operations

TF opens files individually as it copies them, and does not lock disk volumes. Files that are being backed up are opened non-exclusively; files that are being restored are opened exclusively. Files that are being verified on either backup or restore are opened non-exclusively.

You cannot back up or verify a file (in either a backup or restore) that is opened exclusively to another program while the copy is in progress. You cannot replace a duplicate file upon a restore if it is open (exclusively or non-exclusively) to a program. In these cases, TF issues a “File is already open” message and identifies the file. The rest of the backup, restore, or verify is not affected.

It is possible for a program to create a duplicate file while a TF restore is in progress. This duplicate is then purged and replaced with the file from the backup tape.

A program may alter the contents of a file between the time it is copied then opened for verification. In this case, either a “Data doesn’t match source file” error occurs on backup, or an “Update times don’t match. Not Verifying” error occurs on restore.

Using TF with FMGR Files

You can use TF to back up and restore FMGR files, with the limitations listed below. Problems may arise when using TF to access a file if any of the following apply:

- The name contains a “/”
- The name contains more than one “.”
- The name begins or ends with a “.”
- The name ends with “.DIR”
- The name contains a “>” or “[” anywhere after the first character
- The CRN contains a “/”, “.”, “>”, or “[”

TF cannot handle such file names correctly because of the special significance the file system attaches to these characters.

If you try to back up or restore such files (including restoring from FC tapes), you may encounter the following problems:

- The source file descriptor reported in the “Copying ...” message may not agree with the actual source file.
- The destination file may have part of its name dropped.
- The file type may be lost, causing the type to default to 3. For type 1 or 2 files, this usually makes the file data inaccessible. (This can only happen for files on FC tapes with names containing “>” or “[”.)
- The file may be mistaken for a directory file, causing its data to be lost. (This can happen with files whose names end in “/” or “.DIR”.)
- Part of the name may be interpreted as a subdirectory name (as in A/B::C).
- An “illegal name” error may be reported.

TF

Some of these problems can occur without TF detecting an error. If you use TF for both backup and restore with the Verify option, the chance of error detection increases.

For backup operations, verification prevents undetected data loss except for files whose names end in “.DIR”, although file names may be altered without warning. Further, any data loss during backup of files with the “.DIR” extension may cause a verify error on restore, so data loss does not go undetected on both backup and restore.

On a restore, verification prevents undetected data except for possible default to file type 3 (FC tapes only), although file names may be modified or misinterpreted without warning.

- When you use the CO command, you may not use the “@” character to specifically identify file names that contain the “@” character. However, you may copy files whose names contain an “@”.
- There is no incremental backup capability for FMGR files.

Using TF with FC Tapes

In general, you can use TF to read tapes written by the FC utility. This capability may be desirable because TF can restore files to directories, while FC is limited to using FMGR cartridges. To make file conversion convenient, use FC to back up FMGR files, then use TF to restore these files to directories.

When you read FC tapes with TF, the following limitations apply:

- All of the restrictions with regard to special characters in file names and CRNs are applicable (see the preceding section “Using TF with FMGR Files” for details).
- The TF LH command can only list tape headers for TF tapes. When you use LH with an FC tape, it indicates that the tape is in FC format but does not list the header. Use the FC LH command to list the FC tape header.
- You must use the FC CL (Cartridge List) command to obtain a cartridge list. TF does not include a comparable command.
- TF cannot restore sparse files (files with missing extents).
- TF can restore files from multi-tape FC backups, but you must issue a separate copy command for each tape. Individual files that cross FC tapes cannot be restored. For each tape in a multi-tape set, TF reports “FC file error *file*” for every file in the backup set that is not stored on the tape you are currently restoring. Spurious errors such as “Tape runaway” or “tape i/o error” may be reported at the end of each 9-track tape, except for the last tape in the set.
- TF does not include the special FC features provided for FMGR cartridges, such as automatic cartridge allocation and the pre-restore check, to determine whether the cartridge can hold all the files to be restored. (TF creates global directories, rather than allocating cartridges.)

When you use TF to restore FC from tapes, the error message “FC file error *file*” indicates a problem with the file, which may be caused by any of the following:

- A previously reported tape I/O error that, while not directly associated with a particular file, affects the file reported in the message.
- A tape format error (corrupt tape).
- The file reported in the message is a sparse file.
- A disk error occurred when the file was backed up (the error is recorded on the tape).
- The file is continued across tapes.

It may be possible to get more specific information on any error that occurs by restoring the file with FC instead of TF. Improved error recovery may be possible using the FC I option.

Multi-Tape Backup and Restore

When a file cannot fully fit on a tape, the incomplete portion is removed from the tape and the remaining contents are verified (if you selected the Verify option). TF then prompts you to mount the next tape and continues the copy, beginning with the file that did not fit on the previous tape. If a single file is too large to fit on one tape, use a larger tape or back up the file by some other means.

In a multi-tape backup, each tape in the set bears the same title and date in the tape header. The header also includes a tape number to identify the tapes in order, from tape 2 to the end of the set (the first tape in the set is not assigned a number). The title and date from the first tape are always repeated on subsequent tapes in the set, even when new tapes are appended. You can use the LH command to examine the tape header.

When you restore from a multi-tape backup, TF prompts for each tape in succession, restoring and verifying each tape before prompting for the next. TF issues a message if a tape is mounted out of sequence, and it includes the option to remount in the correct sequence or to skip one or more tapes in the set.

UNIX Compatibility

The tape format that TF writes is fully UNIX compatible. However, there are incompatibilities between FMP files and UNIX files that have nothing to do with the tape format itself. TF provides complete or partial solutions for some of these problems, but in some cases, the only course to take is to avoid the problem.

Multi-tape TF backups can be read by TAR, but not in a single command. You must invoke TAR separately for each tape.

Compatibility problems between UNIX files and FMP files primarily involve the format of the data in files; however, there are significant differences in the way that FMP and UNIX name files. Compatibility issues are discussed in the sections that follow.

File Formats on FMP and UNIX

FMP files consist of a series of records, each of which contains a series of bytes (or characters). UNIX files simply consist of a series of bytes. Bytes are not grouped into records on UNIX files.

For text files, FMP puts one line of text in each record. On UNIX files, because there are no records, a special character called “newline” is used to separate lines of text. Note that this

TF

character cannot be contained within a line of text. The UNIX “newline” character is a byte with a decimal value of 10 and is usually referred to as the line-feed character.

If appropriate options are specified, TF converts text files to UNIX format when copying to tape and converts them back when reading tapes. TF converts text files as it copies them by changing the divisions between records into line-feed characters, or vice-versa.

The presence of tabs in text files that originate on a UNIX operating system may require additional conversion. See the suggestions given below under “Summary of Recommendations”.

For our purposes, anything other than a text file is considered a binary file. Binary files do not generally require conversion in order to be copied between FMP and UNIX systems. The contents of binary files can be treated as a series of bytes on either system. FMP uses leading and trailing length words to separate records in variable-record-length files (type 3 and above). TF considers these length words to be part of the file data for binary files. (TF uses forced type 1 access for binary files to bypass the record processing normally done by FMP. Because a UNIX operating system provides no way to indicate records in files, the raw data underlying the FMP record structure is copied directly.)

Porting binary files assumes a common interpretation of the data format between programs on the two types of systems. Differences between hardware representations of numeric data must also be taken into account. Be aware that on some non-HP computers, bytes are stored in inverse order within words and double words, so that bytes may have to be reversed before numeric values are interpreted. If variable-record length binary files are ported to or from a UNIX operating system, the programs that access the files on a UNIX system must understand how FMP uses leading and trailing length words to mark records.

For FMP type 1 and 2 files that originated on UNIX, the size of the main and extents may be larger than the file size on a UNIX system. The end of the file is typically padded with undefined data. This may cause problems unless the file data itself includes information that enables the program that reads the file to determine where the data ends. Type 1 and 2 files that originate in the FMP file system normally contain such information and therefore should not cause problems.

Copying Files Between FMP and UNIX

When files are copied between FMP and a UNIX operating system, text files generally need to be converted as described above, whereas binary files are left unchanged. In order to do this correctly, TF must be able to distinguish text files from binary files. How this is done depends upon the situation, as explained below.

Copying from FMP to UNIX-TAR-Compatible TF Tape to a UNIX Operating System

When TF copies files to tape, it assumes that type 4 files are text and files of all other types are binary. If you specify the X (UNIX) option, type 4 files are converted and other file types are left unchanged. However, if you want text files to remain unconverted, thus leaving them in the FMP record format (with leading and trailing length words), you may omit the X option. If you do not specify X, no files are converted.

Text files copied to a UNIX operating system in this way should be treated as binary files if they are copied back to FMP from a UNIX TAR tape. In other words, you should select the N option. See the detailed information below.

Type 4 files copied to tape with the X option are converted by changing record boundaries into line-feed characters (UNIX “newlines”) so that the files are usable as text files on a UNIX

operating system. If the files contain any line-feed characters before conversion, the original line-feeds will appear the same as the new ones, and both the original and the new ones added by TF are treated as “newlines” on a UNIX operating system. Furthermore, both kinds are changed into record boundaries when read from tape by TF.

The result of copying such a file to tape (with the X option) and back to disk is that the line-feeds are deleted, and the records are split where the line-feeds were. This would seriously corrupt a binary file because binary files are likely to contain line-feeds (bytes with decimal value 10). Therefore, you should not use the X option if any of the type 4 files might contain line-feed characters. If a file that contains line-feeds was converted when copying it to tape, you cannot copy it back from tape correctly, regardless of which options you specify.

In the process of converting type 4 files copied to tape with the X option, records longer than 255 characters are truncated. No warning is given when this is done.

Although files other than type 4 may sometimes contain text, especially type 3 files, such files are not converted because they do not always contain text. Any attempt to convert files that do not contain text causes them to be corrupted simply by copying them to tape and back using TF, as explained above. In contrast, type 4 files almost always contain text, so converting them is unlikely to cause problems.

Copying from UNIX-TAR-Compatible TF Tape to FMP

When copying files from TF tapes, TF can tell which files were converted when copying to tape and therefore automatically knows which files to convert back.

Copying from a UNIX Operating System to TAR Format Tape to FMP

Because no distinction is made on UNIX TAR tapes between text and binary files, TF cannot automatically tell which files need to be converted when reading TAR tapes. Because text files are the most typical, TF assumes all files on TAR tapes are text files and converts them accordingly. You can use the N option, which inhibits conversion, to override this default so that binary files can be copied. (Applying the conversion used for text files to a binary file causes the data to be copied incorrectly.)

When TF converts files that are being copied from a TAR tape, lines are split into separate lines after every 256 characters. When this record splitting occurs, it will cause verification errors.

UNIX TAR tapes have no information to indicate the proper FMP file type. When TF copies from TAR tapes, it assumes type 4 unless you specify another file type in the destination parameter. For type 2 files, you must also specify the proper record length. If you do not specify the appropriate file type (and record length, if applicable), FMP may not be able to access the file correctly.

Normally, you should copy files that originated on FMP back to their original file type. File type 1 is usually appropriate for binary files that originated on a UNIX operating system. An understanding of FMP file types may be necessary to determine what file type is appropriate in a given situation. (Refer to the *RTE-A Programmer's Reference Manual*, part number 92077-90007, for a description of file types.)

To restore files from a UNIX TAR tape that contains mixtures of binary and text files, or mixtures of various file types, you must specify the N option for some files and not for others, and you must specify different file types for different files. You may use the GR command to do this, by

TF

specifying different destinations and options for each command in the group. Using a standard naming convention to distinguish different kinds of files facilitates this, as shown in the following CO command group:

```
TF: gr
TF: co 8{@.txt}
TF: co 8{@.bin} :::1 n
TF: eq
```

This command uses the default action for files matching `@.txt`, treating them as text files, converting them accordingly, and copying them into type 4 files. However, files matching `@.bin` are treated as binary files and not converted (due to the `N` option), and are copied into type 1 files (as specified in the destination parameter).

The above example assumes that files on the tape were in the working directory and were named relative to the working directory when they were copied to tape. Otherwise, you must specify a directory or specify the `E` qualifier (`@.txt.e`) to indicate that all directories are to be searched.

If you do not use a standard naming convention, you must expand the group to name each file individually. However, if a tape contains files that are mostly of one type, with a relatively small number of files that are exceptions, it is possible to name only the exceptions at the beginning of the group and then end the group with a single blanket command to cover all the other files. For example:

```
TF: gr
TF: co 8{firsttype1 secondtype1} :::1 n
TF: co 8{firsttype2 secondtype2 thirdtype2} :::2:200 n
TF: co 8{fourthtype2 fifthtype2} :::2:200 n
TF: co 8
TF: eg
```

This command works because when files are copied from tape, only the first of the matching commands is used, even though a file may match the source mask of more than one CO command.

This means the destination and options used for a particular file are the ones specified in the first command with a matching source mask. Thus, files selected by each CO command in a group are automatically excluded from subsequent commands in the group.

In the grouped command shown above, each file specifically named in the first three commands in the group is copied with the `N` option and the specified file type of 1 or 2. All the other files are copied into type 4 files (the default) with no options. The final command specifies no mask and thus matches all the files not matched by the previous commands.

You can also use the `DE` command effectively to save typing long lists of files.

Summary of Copying Recommendations

Copying Files to Tape with TF

In general, you should use the X option when you copy files to tape with TF if the tape is to be read by the UNIX TAR utility. The X option causes type 4 files to be converted to the UNIX text file format.

For files that are to be copied to a UNIX operating system by using TF with the X option, always keep text in type 4 files, and never put binary data in type 4 files. That way TF knows which files to convert.

Do not use the X option with type 4 files that may contain line-feed characters because this may cause the files to be corrupted by the conversion processes used in going to tape and back again. Type 4 files are usually text files, which do not usually contain line-feeds, because they are control characters (control-J). Therefore, this is only a problem if a type 4 file contains binary data (likely to contain line-feed characters, which are simply a byte with decimal value 10) rather than text or if it is a text file that actually contains a control-J.

When type 4 files are copied to tape with the X option specified, records longer than 255 characters are truncated without warning.

You may omit the X option when you move files to a UNIX operating system if you want to leave text files in the FMP record format.

Copying Files from Tape with TF

When you copy text files from UNIX TAR tapes, note that text files from a UNIX operating system frequently contain tabs, but RTE does not support the use of tabs in text files. Files that contain tabs cannot be compiled or printed. However, you may use EDIT/1000 to expand tabs into spaces. Use the EDIT/1000 T command to set tab stops and turn screen mode display functions off (SE DF OF). Then go through the entire file in screen mode, reading each screen and going on to the next using the control-N command (control-Q after the last screen).

When you copy from a UNIX TAR tape, you must specify the N option for binary files to inhibit the conversion process, which is only needed for text files. Specifying the N option for text files, or omitting it for binary files, causes incorrect data in the resulting files. For tapes that contain a mixture of text and binary files, you may specify N for some files but not others by using grouped CO commands.

Text files stored on a UNIX operating system in FMP format are like binary files and need not be converted when they are copied back to FMP. Therefore, you should specify the N option for these files as well. (Such files may result from copying text files to a UNIX operating system without specifying the X option.)

When you copy files from a UNIX TAR tape, the file type defaults to 4. If the files are not to be copied into type 4 files, you must specify the appropriate file type in the destination parameter, to ensure that the files are accessed correctly by FMP. For type 2 files, you must also specify the record length. If several file types are mixed on a single tape, you may use a grouped CO command to specify the different types.

TF

When you copy files from a UNIX TAR tape without using the N option, lines are split into separate lines after every 256 characters. No warning is given when this is done.

When you copy binary files, whether from FMP to a UNIX operating system or from UNIX to FMP, compatibility of the data format must not be taken for granted. The various issues concerning binary files discussed in the section “File Formats on FMP and UNIX” in this chapter should be considered.

Directory File Names

All FMP directories have a type extension of .DIR, whereas a UNIX operating system imposes no such restriction and generally uses ordinary names (with no periods) for directories. Therefore, TF removes the .DIR from names of directory files when copying to tape and appends .DIR to directory file names when reading from tape.

Thus, the files EXAMPLE.DIR and EXAMPLE are both called EXAMPLE when read from a TF tape to a UNIX system, which causes problems.

File Name Letter Case

A UNIX operating system allows both uppercase and lowercase letters in file names and treats the two cases differently. FMP allows you to specify names in either case, but uses only uppercase on disk. Thus, the file names ‘myfile’, ‘MYFILE’, and ‘MyFile’ refer to three distinct files on a UNIX system; but on RTE, all three refer to the same file.

When you copy from a UNIX TAR tape to an FMP disk, all file names are converted to uppercase. However, if the tape contains multiple files that have the same name except for letter case, this conversion causes duplicate names, and thus, only the last of the multiple files is restored to FMP. The message “Replacing *filename* (created during this copy command)” displays for each duplicate file name.

Because file names in a UNIX system are generally lowercase, TF converts all file names to lowercase when copying to tape. This is transparent to RTE because TF converts file names to uppercase when reading tapes. However, copying files from a UNIX operating system to FMP and back to UNIX can cause problems because the effect in this case is to convert all file names to lowercase.

Periods Used in File Names

FMP file names use a period (.) to separate the file name from the type extension. No other use of a period is allowed. A UNIX operating system, however, treats the period as any character in a file name. When TF encounters a period in an unexpected position in a file name or path name, it converts the period to an underscore (_). This includes any of the following placements:

- A period at the beginning of a name
- A period at the end of a name
- A period in a directory name
- A period that precedes other periods (for example, the first period in file a.b.c)

The following shows the resulting UNIX path names when periods are converted to underscores and the implicit .DIR type extension for directory files is inserted. Note that TF does not recognize the special names “.” and “..” used by a UNIX operating system.

Before	After
/aaa/bbb.ccc	aaa/bbb.ccc
/aaa/bbb.ccc (directory file)	aaa/bbb_ccc.dir
/aaa/bbb.ccc/ddd	aaa/bbb_ccc/ddd
aaa/bbb/.ccc	aaa/bbb/_ccc
aaa/bbb/ccc.	aaa/bbb/ccc_
aaa/bbb.ccc.ddd	aaa/bbb_ccc.ddd
aaa/bbb.ccc.ddd (directory file)	aaa/bbb_ccc_ddd.dir
.aaa	_aaa
./aaa	_/aaa
bbb/.. (directory file)	bbb/____.dir
bbb/./aaa	bbb/____/aaa

UNIX users who execute “tar cv .” get files on the TAR tape whose names begin with “./”. If you restore such tapes with TF, the command

```
tf co tape
```

creates a directory called “_” in the working directory and copies all the files into that directory (TF changes “./example” to “/example”). To avoid this problem, use “tar cv *” rather than “tar cv .”, or use the following command instead of the one above:

```
tf co tape{_/@.}@ @.@"
```

Special Characters

File names that contain blanks or the characters “[”, “>”, and “:” cannot be handled correctly by FMP because they are used as delimiters.

The “@” character can cause problems in some contexts because it is the FMP wildcard character. (In a UNIX operating system, the “*” is the wildcard character.)

File Name Length—Moving Files from FMP to UNIX

FMP files can have up to 21 characters: 16-character names and a period followed by a 4-character type extension. UNIX file names are limited to a total of 14 characters. Therefore, FMP files that have more than 14 characters (including the name, period, and type extension) cannot be moved to a UNIX operating system. What happens if you attempt to do so is a function of the TAR utility. (Some newer UNIX versions do allow file names greater than 14 characters.)

File Name Length—Moving Files from UNIX to FMP

UNIX file names may contain a period followed by more than four characters. Because FMP interprets the last period as the type extension delimiter for the file name, UNIX file names are truncated to four characters following the last period. No warning is given when this occurs.

TAR allows path names of up to 99 characters; FMP allows only 63 characters in a path name. If a path name from a TAR tape is longer than 63 characters (after any required conversion), TF returns the message “Name too long. Skipping *filename*”.

TF

Linked Files

Linked files are not supported by FMP, and cannot be restored.

Appending to Tapes

Tapes created by one utility should not be appended using the other utility, because TF and TAR do not append to tape in the same way.

File Types, Security Codes

A UNIX operating system does not recognize file types or security codes, which are lost when FMP files are moved to a UNIX system. Files moved from UNIX to FMP default to type 4, security code 0.

Time Stamps

UNIX TAR tapes do not have a create time for each file. Refer to the earlier section on “Time Stamps” for details.

Root Directory

The FMP root directory may only contain directory files; the UNIX root directory may also contain ordinary files. You can restore ordinary files from the root directory by specifying a directory other than the root in the destination parameter.

When you select these files from a UNIX TAR tape, you must specify the type extension or use a wildcard. Otherwise, TF assumes a DIR type extension and just selects the directory files from the root directory.

Header and Final Checksum Dummy Files

TF includes a header file at the beginning of each tape and a dummy checksum file after the last file on the tape. These are special files that are not restored. To TAR, however, these are ordinary files and they are extracted with the rest of the files.

The tape title is used as the path name for the header file. If you do not specify a title for the tape header, the default title (the CO command itself) will contain embedded spaces and could contain slashes. This could result in extraneous files being created by TAR.

The checksum file always has the path name “.CHECKSUM”.

Because neither the header nor the checksum file contains data, the use of extra disk space is not a problem.

TF Tape Format

The TF tape format and header basics are shown below.

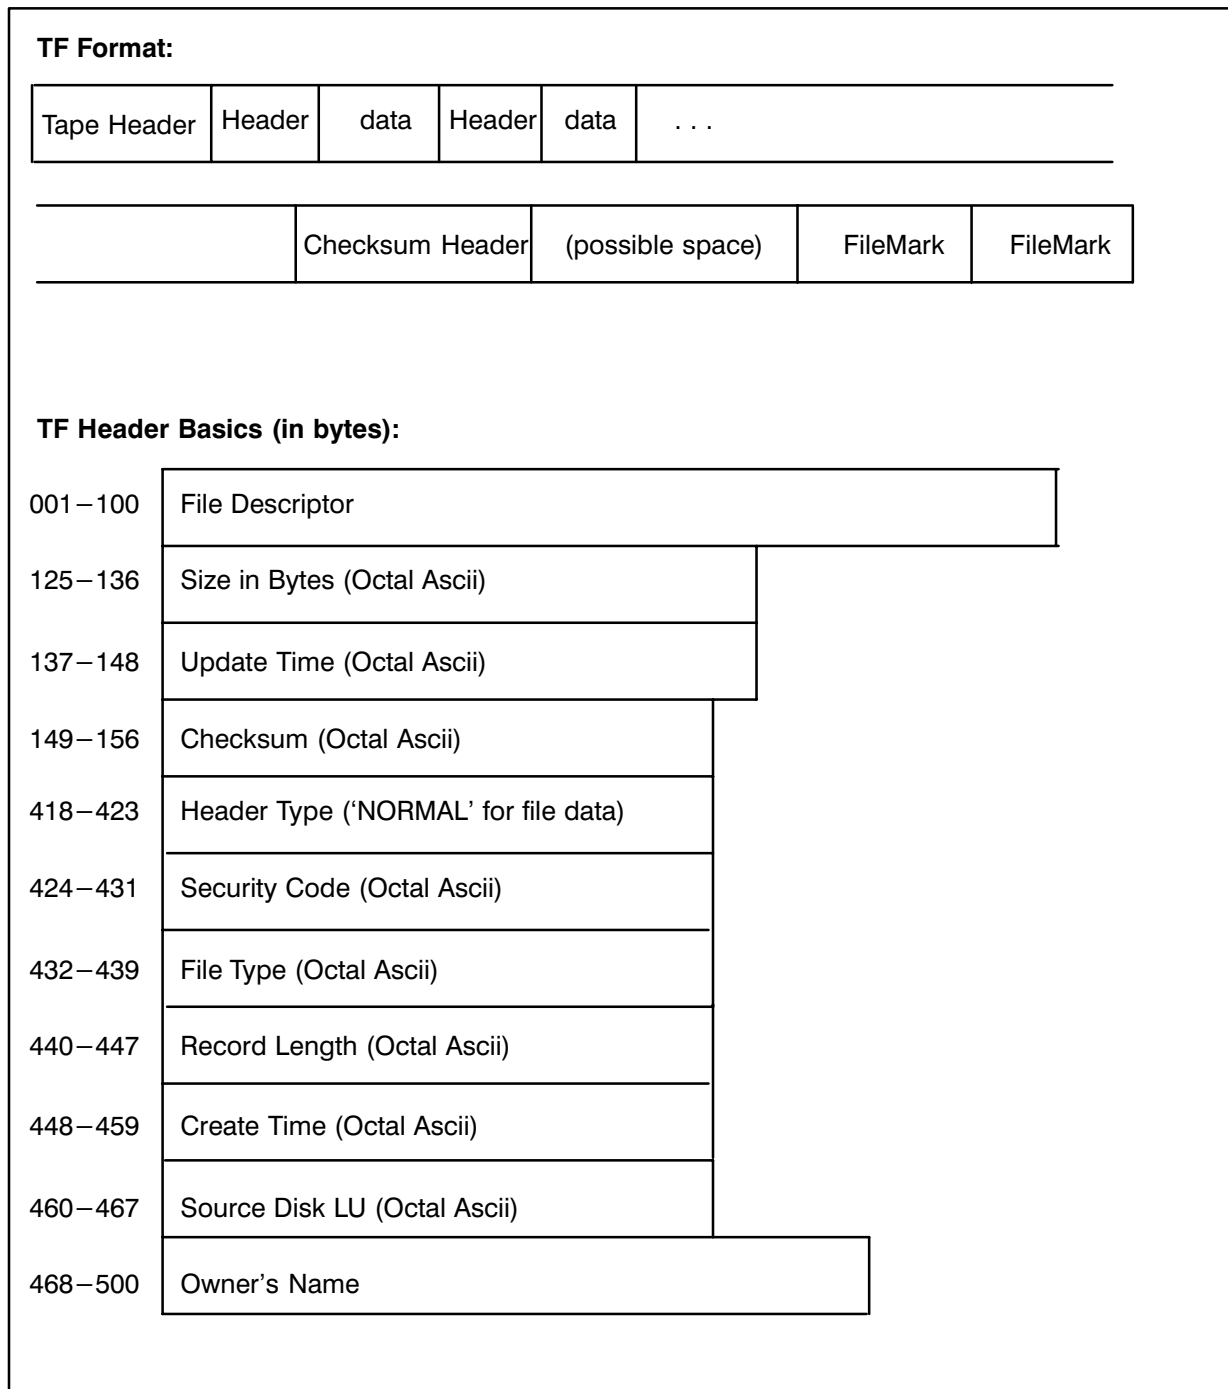


Figure 2-7. TF Tape Format and Header Basics

TF

Installing TF

Load the TF utility using the LINK command file #TF.

The global directory /SCRATCH should be created during system installation and put on an LU with a lot of free space. This ensures trouble-free operation of TF and other system utilities that use the /SCRATCH directory for scratch files.

The file >TF000 must be located in one of two places. If the global directory /CATALOGS exists, place >TF000 within it. If /CATALOGS does not exist, you can place >TF000 on an FMGR system cartridge; however, if global directory /CATALOGS is later created, you must move >TF000 to /CATALOGS for TF to run.

You may want to increase TF's EMA size when you install TF. EMA size should be increased if you want to use the TF GR command to group more than a few CO commands. If the EMA size is not large enough for all the CO commands in a group, TF displays the following message and exits:

```
EMA size of TF not large enough for a group of this size.
```

Each additional page of EMA lets you include at least 20, and possibly as many as 100, more CO commands in a group, depending upon the complexity of the individual commands.

You can increase the EMA size when TF is loaded by increasing the number of pages specified in the EM command in the LINK command file, #TF. Alternatively, you can use LINK to relink an existing TF program file, as:

```
CI> link
link: lk tf::programs
link: em 30
link: en
```

Index

A

ARSTR physical restore utility, 2-2
 ABort, 2-3
 break detection, 2-12
 calling ARSTR, 2-2
 checking the disks, 2-9
 command examples, 2-8
 commands, 2-2
 disk locking, 2-10
 ENd, 2-3
 error messages, 2-14
 example, 2-13
 EXit, 2-3
 HElP, 2-4
 LH List Header, 2-4
 LL List Device, 2-4
 loading ARSTR, 2-12
 offline system, 2-12
 order of disk restoration, 2-10
 RE command options, 2-6
 REstore, 2-5
 restore operation, 2-9
 restoring the data, 2-11
 RW Rewind, 2-7
 save definition records, 2-10
 sectors per track must match, 2-9
 TApe LU, 2-7
 total number of tracks may differ, 2-9
 UE User Error Handling, 2-8
 unit restore, 2-10
 verification, 2-11

ASAVE physical backup utility, 2-19
 ABort, 2-21
 break detection, 2-29
 calling ASAVE, 2-19
 command options, 2-26
 commands, 2-20
 data record, 2-31
 disk locking, 2-26
 ENd, 2-21
 end-of-data record, 2-32
 error messages, 2-33
 example, 2-32
 EXit, 2-21
 header records, 2-30
 HElP, 2-21
 LH List Header, 2-21
 LL List Device, 2-22
 loading ASAVE, 2-29
 RW Rewind, 2-22
 SA command options, 2-23
 SAve, 2-22
 save definition record, 2-31

 save file records, 2-27
 save operation, 2-26
 saving disks to tape, 2-27
 tape and file formats, 2-29
 TApe LU, 2-24
 tape positioning, 2-27
 tape records, 2-30
 TItle, 2-24
 UE User Error Handling, 2-25
 unit save LU checking, 2-26
 verification, 2-28

B

backing up
 disk LUs, 2-19
 temporary files, 2-233

backup
 logical, 1-2
 online, 1-2
 utilities
 ARSTR, 2-2
 ASAVE, 2-19
 COPYL, 2-34
 DRSTR, 2-36
 DSAVE, 2-52
 FC, 2-85
 FST, 2-172
 LIF, 2-219
 overview, 1-1
 TF, 2-230

C

caching scheme, 1-9
cartridge, tape drive, CS/80 definition, 1-9
COMND program, 2-12
continuous streaming mode, 1-9
COPYL. *See* disk-to-disk copy utility
CS/80, cartridge tape drive, definitions, 1-9
CS/80 disk formatting utility (FORMC), 2-137
 ABort, ENd, and EXit commands, 2-139
 break detection, 2-138
 calling FORMC, 2-137
 command execution, 2-138
 converting spares during tape formatting, 2-145
 data loss during sparing, 2-147
 default record interleave factor, 2-142
 device driver status, 2-138
 disk formatting information, 2-141
 disk formatting process, 2-143
 dismounting LUs, 2-142
 error messages, 2-151
 Format command, 2-140

- formatting a system disk, 2-143
 - formatting operation, 2-141
 - FORMC commands, 2-139
 - help command, 2-140
 - inserting the cartridge, 2-144
 - loading FORMC, 2-151
 - re-certification and spare conversion, 2-144
 - readying a tape for verification, 2-148
 - SPare command, 2-140
 - sparing defective blocks, 2-146
 - sparing operation, 2-146
 - specifying a start track or block for verification, 2-149
 - specifying an interleave factor, 2-142
 - specifying the number of tracks or blocks to verify, 2-149
 - specifying the track when sparing, 2-146
 - tape certifying and skip sparing, 2-144
 - tape formatting information, 2-144
 - tape formatting process, 2-145
 - verification process, 2-150
 - VERify command, 2-141
 - verifying an entire disk LU or tape, 2-148
 - verifying operation, 2-148
- D**
- disk, caching, 1-9
 - disk blocks, 1-5
 - address, 1-5
 - CS/80 definition, 1-9
 - postamble, 1-5
 - preamble, 1-5
 - read/processing time calculated for interleave factor, 1-8
 - sparing defective ones, 2-146
 - status bits, 1-7
 - testing, 1-5
 - user data area, 1-5
 - disk formatting utility (FORMF), 2-152
 - ABort, ENd, EXit commands, 2-154
 - break detection, 2-153
 - calling FORMF, 2-152
 - command execution, 2-153
 - device driver status, 2-153
 - error messages, 2-158
 - FORmat command, 2-155
 - formatting operation, 2-156
 - formatting process, 2-157
 - FORMF commands, 2-154
 - HElp command, 2-155
 - specifying an interleave factor, 2-156
 - supported disk drives, 2-152
 - VERify command, 2-155
 - verifying operation, 2-157
 - disk-to-disk copy utility (COPYL), 1-3, 2-34
 - calling COPYL, 2-34
 - copy operation, 2-35
 - error messages, 2-35
 - DRSTR physical restore utility, 2-36
 - ABort, 2-37
 - break detection, 2-46
 - calling DRSTR, 2-36
 - checking the disks, 2-42
 - commands, 2-36
 - disk locking, 2-43
 - ENd, 2-37
 - error messages, 2-47
 - EXit, 2-37
 - file-by-file (LU) restore, 2-39
 - header records, 2-45
 - HElp, 2-37
 - LH List Header, 2-38
 - LL List Device, 2-38
 - loading DRSTR, 2-46
 - MEdia, 2-39
 - order of disk restoration, 2-43
 - RE command options, 2-41
 - REstore, 2-39
 - restore operation, 2-42
 - restoring fixed disks from removable disks, 2-44
 - restoring the data, 2-45
 - sectors per track must match, 2-42
 - total number of tracks may differ, 2-42
 - UE User Error handling, 2-41
 - UNit restore, 2-41
 - unit restore, 2-40, 2-44
 - verification, 2-46
 - VERify data, 2-41
 - DSAVE physical backup utility, 2-52
 - ABort, 2-54
 - break detection, 2-62
 - calling DSAVE, 2-52
 - commands, 2-53
 - disk locking, 2-60
 - ENd, 2-54
 - error messages, 2-66
 - EXit, 2-54
 - header records, 2-64
 - HElp, 2-54
 - LH List Header, 2-54
 - LL List Device, 2-56
 - loading DSAVE, 2-63
 - MEdia, 2-56
 - no locking, 2-58
 - removable disk media and file formats, 2-64
 - removable disk overwriting, 2-60
 - SA command options, 2-57
 - SAve, 2-57
 - save file records, 2-61
 - save operation, 2-59
 - saving to the removable disk, 2-61
 - TItle, 2-58
 - UE User Error handling, 2-59
 - UNit save, 2-58
 - unit save LU checking, 2-60
 - verification, 2-62
 - VERify, 2-58

E

Error Rate Test utility (ERTSH), 1-5, 2-67, 2-129, 2-133
BP Build Test Pattern command, 2-69
break detection, 2-78
commands, 2-68
CY Define Cylinder Range command, 2-69
error interpretation, 2-80
error messages, 2-83
ERTSH structure, 2-67
EXit command, 2-70
HD Define Disk Head Range command, 2-70
help command, 2-69
initializing and testing a new disk, 2-75
installation procedure summary, 2-79
LL Extra Log Device command, 2-70
LU Selecting Disk LU to Test command, 2-71
P# Set Test Pass Count command, 2-71
RO Execute Read-Only Error Rate Test command, 2-72
SC Define Disk Sector Range command, 2-72
specialized use of ERTSH, 2-78
SS Spare Sectors command, 2-73
testing disks formatted with FORMF, 2-75
testing the functionality of a disk and its controller, 2-67
TP Test Pattern command, 2-73
typical use of ERTSH, 2-75
using ERTSH, 2-75
WR Write/Read Error Rate Test command, 2-74
ERTSH. *See* Error Rate Test utility
extent, file, 2-190
eliminate using FC, 2-94
reducing for performance, 2-107

F

FC. *See* File Copy
file
backup utilities, 1-2
extent, 2-94
interchange utilities, 1-4
FC, 2-85
FST, 2-172
LIF, 2-219
overview, 1-1
TF, 2-230
File Copy (FC), 1-4, 2-85
ABort command, 2-89
brief, full status display format, 2-93
calling FC, 2-85
cartridge lock, open, 2-95, 2-108
CF name comment file command, 2-89
CL Cartridge List command, 2-89
clear destination disk, 2-94
CO command examples, 2-96
CO command options, 2-93

CO command source and destination parameters, 2-92
command summary function (?), 2-88
COpy command, 2-90
copy single volume of multi-volume tape set, 2-95
DEfault command, 2-99
destination disk handling, 2-106
display required tape length, 2-95
DL Directory List command, 2-100
ECho command, 2-102
eliminate extents, 2-94
error handling in transfer files, 2-110
error messages, 2-112
EXit command, 2-102
extents, file, 2-94, 2-107
FC commands, 2-86
globals used in transfer files, 2-108
GRoup CO commands, 2-102
ignore data errors, 2-94
keep tape online, 2-95, 2-103, 2-106
LC List Comment files command, 2-103
LH List Header files command, 2-103
LL List device command, 2-103
loading FC, 2-108
master security code, 2-91
performance considerations, file copy operations, 2-107
purge source file, 2-95
recover unused space, 2-96
replace duplicate files, 2-94
SChatch area definition command, 2-104
SKip volume option in multi-volume read, 2-105
tape directory list format, 2-101
tape handling, 2-105
TITle command, 2-104
TRansfer command (to/from command file), 2-104
verify transferred data integrity, 2-96
File Storage to Tape utility (FST), 2-172
append option (A), 2-188
appending data, 2-201
backing up using file masking, 2-193
Backup Bits, 2-189
BAckup command, 2-176
preserving FMGR security codes, 2-176
brief option (B), 2-189
building a new directory file, 2-204
calling FST, 2-173
clear option (C), 2-189
command options summary, 2-188
command stack, 2-176
commands, 2-175
consecutive backups, 2-201
D, K, N, and S qualifiers, 2-193
delta backups, 2-197
DF Directory File command, 2-177
disk directory file, 2-205
DL List Directory command, 2-177
duplicate option (D), 2-189

- EM SHEMA Size command, 2-178
- end-of-file position ignored with W option, 2-192
- error handling, 2-209
- error messages and warnings, 2-210
- EXit command, 2-178
- extents, file, 2-190, 2-206
- faulty option (F), 2-189
- file masking and renaming, 2-193
- FST.RC start-up file, 2-173, 2-174
- full backups, 2-197
- GO begin backup/restore command, 2-178
- HElP command, 2-178
- incremental backup, 2-197
- inhibit option (I), 2-189
- installing FST, 2-209
- keep option (K), 2-190, 2-201
- LC List Comment file command, 2-179
- LH List Header command, 2-179
- LI List selected files sommand, 2-179
- LL select log device/file command, 2-179
- LN List Non-selected files command, 2-180
- lock option (L), 2-190
- mindir option (M), 2-190
- MT specify tape LU command, 2-180
- multiple reels, 2-201
- NEXt command, 2-181
- normal option (N), 2-190
- options, 2-187
- original option (O), 2-190
- POsition command, 2-181
- PREvious command, 2-181
- purge option (P), 2-191
- quiet option (Q), 2-191
- recommended system usage, 2-206
- replacing reserved characters in FMGR file names, 2-205
- REstore command, 2-182
- restoring files from overwritten tape, 2-189, 2-204
- restoring from incremental backups, 2-199
- restoring using file masking, 2-195
- RUn command, 2-183
- rwndoff option (R), 2-191
- SC Select Comment file command, 2-183
- SD Set tape Density command, 2-183
- SEcure command, 2-184
- Shareable EMA (SHEMA), 2-205
- SHow user selected states command, 2-184
- srchapp option (S), 2-191
- streaming during verify pass, 2-192
- streaming mode, 2-182, 2-192, 2-206
- tape format, 2-207
- tape loading, 2-202
- tape positioning on overwritten tapes, 2-204
- TAr command, 2-184
- TF compatibility, 2-202
- TItle command, 2-185
- TRansfer to command file command, 2-185
- UNIX TAR Format, 2-184

- UNselect command, 2-186
- update option (U), 2-191, 2-200
- verify option (V), 2-192
- whole option (W), 2-192
- yes option (Y), 2-192
- Z option, 2-192
- fill number, 1-8
- FORMA. *See* initialize and spare utility
- formatting
 - general information, 1-5
 - utilities, 1-5
 - CS/80 disk (FORMC), 1-5, 2-137
 - disk (FORMF), 1-5, 2-152
 - error rate test (ERTSH), 1-5, 2-67
 - initialize and spare (FORMA), 1-5, 2-129
 - offline disk formatting (FORMT), 1-5, 2-160
- FORMC. *See* CS/80 disk formatting utility
- FORMF. *See* disk formatting utility
- FORMT. *See* offline disk formatting and initialization
- FST. *See* file storage to tape utility

H

- HP computer systems file copy utility (LIF), 2-219
 - Calling LIF, 2-220
 - CO command, 2-222
 - DL Directory List command, 2-223
 - EXit command, 2-224
 - HElP command, 2-224
 - INitialize command, 2-224
 - LIF commands, 2-220
 - LIF error handling, 2-229
 - LlSt command, 2-225
 - LL set Logical List device command, 2-225
 - MC Mount Cartridge command, 2-225
 - naming conventions, 2-219
 - PK Pack Cartridge command, 2-226
 - PURge command, 2-226
 - RN rename command, 2-226
 - STore command, 2-227
 - SV severity command, 2-228
 - TR TRansfer control command, 2-228

I

- initialize and spare utility (FORMA), 2-129
 - byte/index address, 2-132
 - calling FORMA, 2-131
 - converting from FORMF to FORMA, 2-133
 - definitions, 2-129
 - disk error decoding, 2-136
 - error handling, 2-134
 - error messages, 2-134
 - ERTSH works with FORMA, 2-129, 2-133
 - FORMA commands, 2-131
 - help command, 2-131
 - IM Initialize Media command, 2-131
 - installing a new internal hard disk, 2-133
 - SS command options, 2-133

- SS Spare Sector command, 2-132
 - typical track format, 2-130
- initializing
 - and testing a new disk, 2-75
 - hard disks, 1-7
- interleave factor, calculation, 1-8

L

- LIF. *See* HP Computer Systems File Copy
- logical, backup, 1-2

M

- memory-based system, ASAVE and ARSTR, 2-12

O

- offline
 - restore, 1-3
 - system, ASAVE and ARSTR, 2-12
- offline disk formatting and initialization utility (FORMT), 2-160
 - calling FORMT, 2-160
 - ENd command, 2-161
 - error messages, 2-170
 - FOrmat command, 2-161
 - formatting operation, 2-164
 - formatting process, 2-165
 - FORMT commands, 2-161
 - FORMT functions, 2-160
 - generating a FORMT system, 2-160, 2-163
 - help command, 2-162
 - INitialize command, 2-162
 - initializing operation, 2-166
 - REformat command, 2-162
 - reformatting operation, 2-167
 - sector interleaving, 2-165
 - SPare command, 2-163
 - sparing operation, 2-168
 - VERify command, 2-163
 - verify operation, 2-169
- online, backup and restore from/to disks, 1-2

P

- physical, backup utilities
 - ARSTR, 1-3, 2-2
 - ASAVE, 1-3, 2-19
 - compatibility with others, 1-3
 - COPYL, 1-3, 2-34
 - DRSTR, 1-3, 2-36
 - DSAVE, 1-3, 2-52
- preamble, 1-7

S

- SHEMA, size used by FST, 2-178
- skipping, versus sparing, 1-8

- sparing
 - hard disks, 1-7
 - versus skipping, 1-8
- streaming mode, 1-9

T

- Tape Filer (TF), 1-4, 2-230
 - access time for tape files, 2-261
 - alternatives to standard incremental backup, 2-266
 - B qualifier in incremental backup, 2-262
 - backup bits, 2-262, 2-264, 2-265
 - C option used in incremental backup, 2-262
 - calling TF, 2-231
 - CO command, 2-232
 - CO command options, 2-234
 - append to tape (A), 2-235
 - brief logging mode (B), 2-235
 - clear backup bit (C), 2-235
 - ignore errors and file masks (I), 2-236
 - inhibit UNIX to FMP text file conversion (N), 2-236
 - keep tape online (K), 2-236
 - replace duplicate files (D), 2-236
 - Unix compatibility (X), 2-236
 - update (U), 2-236
 - verify files copied (V), 2-236
 - yes option (Y), 2-237
 - CO command source and destination parameters, 2-233
 - copy example without subdirectories, 2-237
 - copy examples using DS, 2-250
 - copy examples with subdirectories, 2-243
 - copying files between FMP/UNIX, 2-274
 - create time (tape), 2-261
 - DEfault command, 2-253
 - delta backups, 2-262
 - directory creation and restore, 2-267
 - directory file names (FMP/UNIX), 2-278
 - disk full errors, 2-270
 - DL Directory List command, 2-254
 - duplicate files during restoring incremental backups, 2-264, 2-266
 - EXit command, 2-257
 - file access during backup/restore, 2-271
 - file extents, 2-233
 - file formats on FMP and UNIX, 2-273
 - file properties, saving and restoring, 2-268
 - file/tape compatibility, FMP/UNIX, 2-273, 2-282
 - full backup, 2-262
 - GRoup copy commands, 2-257
 - HElp command, 2-258
 - incremental backup, 2-262
 - standard procedure, 2-263
 - installing TF, 2-282
 - keep tape online, 2-254, 2-257, 2-258, 2-260
 - LH List Header file, 2-258
 - LL list device command, 2-258
 - maintaining the system time, 2-261

- missing extents, 2-272
- missing time stamps, 2-261
- multi-tape backup/restore, 2-273
- multiple copies of the same backup, 2-264
- relation of DL command to CO command, 2-256
- restoring incremental backups, 2-264
- restoring older versions from incremental backup, 2-265
- sparse files, 2-233, 2-272
- system backup and restore, 2-269
- tape protection and the K (keep) option, 2-260
- temporary files, 2-233
- TF commands, 2-231
- TF tape format, 2-281
- time stamps, 2-261
- TItle command, 2-259
- TRansfer command, 2-260
- UNIX compatibility, 2-273
- update time (tape), 2-261
- using TF with FC tapes, 2-272
- using TF with FMGR files, 2-271
- temporary files, backing up, 2-233
- track size of disks, 1-3
- tracks, sparing and skipping, 1-8