



DMF Recovery and Troubleshooting
Guide for SGI[®] InfiniteStorage

007-3682-011

COPYRIGHT

© 1998, 2000, 2002, 2003, 2004, 2005 Silicon Graphics, Inc. All Rights Reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of Silicon Graphics, Inc.

LIMITED RIGHTS LEGEND

The software described in this document is "commercial computer software" provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

Silicon Graphics, SGI, the SGI logo, IRIX, and XFS are registered trademarks and of Silicon Graphics, Inc., in the United States and/or other countries worldwide.

IBM is a trademark of International Business Machines Corporation. Linux is a registered trademark of Linus Torvalds. UNIX is a registered trademark and X Window System is a trademark of the Open Group in the United States and other countries. All other trademarks mentioned herein are the property of their respective owners.

New Features

For information on changes in DMF functionality, including bugs fixed in recent releases, refer to the **Dependencies** and **News** buttons on the DMF installation interface (`dmmain(8)`).

Record of Revision

Version	Description
2.6.1	January 1998 Original printing. Supports Data Migration Facility (DMF) 2.6.1 running under SGI IRIX systems.
2.6.2	July 1998 Supports DMF 2.6.2 running under SGI IRIX systems.
2.6.2.2	December 1998 Supports DMF 2.6.2.2 running under SGI IRIX systems. The only change associated with this manual was a change in the title.
004	May 2000 Supports DMF 2.6.3 running under SGI IRIX systems. Only minor editing changes are included.
005	October 2000 Supports DMF 2.6.3.2 running under SGI IRIX systems. Only minor editing changes are included.
006	April 2002 Supports DMF 2.7 running under SGI IRIX systems.
007	March 2003 Supports DMF 2.8 running under SGI IRIX and Linux systems.
008	October 2003 Supports DMF 3.0 running under SGI IRIX and Linux systems.
009	August 2004 Supports DMF 3.1 running under SGI IRIX and Linux systems.
010	February 2005 Supports DMF 3.2 running under SGI IRIX and Linux systems.
011	April 2005 Supports DMF 3.3 running under SGI IRIX and Linux systems.

Contents

About This Guide	xv
Related Publications	xv
Obtaining Publications	xvi
Conventions	xvi
Reader Comments	xvii
1. Introduction	1
The Need for <code>dmaudit</code>	2
When to Use <code>dmaudit</code>	3
Disk Cache Manager	3
2. Daemon Internals	5
BFIDs and BFID Sets	5
Migration Life-Cycle	6
Regular Files	8
Migrating Files	9
Dual-State Files	9
Offline Files	9
Partial-State Files	9
Voided-State Files	10
DMF Inode-Resident Data	10
The Bit File Identifier (BFID)	10
DMF File State Field	11
DMF Daemon Database Contents	12
Errors During Migration	15
Remigrating a File	16
007-3682-011	vii

Removing or Modifying a File	16
DMF BFID State Summary	17
3. How <code>dmaudit</code> Detects Errors	19
4. Using <code>dmaudit</code> Interactively	21
Environment Variables	21
Page-Wait Mechanism	21
5. Setting the Initial Configuration	23
Providing a Working Directory	23
Specifying Filesystems to Scan	24
6. Detecting Discrepancies	29
Snapshot Resource Requirements	29
Taking a Snapshot Interactively	30
The snapshot Option	30
The inspect Option	32
The report Option	32
The verifymsp Option	33
The free Option	33
Example of a Report with Discrepancies	34
Taking a Snapshot in Batch Mode	35
7. Examining and Correcting Discrepancies	37
Overview of the Correction Process	37
Selecting Which Errors to Correct	38
Correcting the Errors	39
The INSPECT Menu	39

The <code>bfid</code> Option	41
The <code>search</code> Option	44
Fixing BFID Sets for Which No User Files Exist	45
The <code>accept</code> Option	47
The <code>examine</code> Option	47
The <code>dump</code> Option	49
Returning to the Inspect Menu	49
Fixing Files with Correctable Errors	50
The <code>opt_full</code> and <code>opt_part</code> Options	52
The <code>search</code> Option	53
The <code>nlist</code> Option	54
Returning to the Inspect Menu	55
Cleaning Up Files with Unrecoverable Errors	56
Unrecoverable Errors for Offline Files	56
The <code>remove</code> Option for Offline Files	59
Unrecoverable Errors for Partial-State Files	60
The <code>remove</code> option for Partial-State Files	62
Cleaning Up Files with the Same BFID but Different Sizes	63
Cleaning Up Multiple Files with the Same BFID and Size	67
The <code>edit</code> Option	70
The <code>duplicate</code> Option	72
The <code>transfer</code> , <code>remove</code> , and <code>nondup</code> Options	74
BFID Sets That Cannot Be Immediately Corrected	78
The <code>apply</code> Option	80
8. Changing the Configuration	83
The <code>workdir</code> Option	84

The fileysys Option	85
The invalid Option	86
The dmfcconf Option	88
The remove Option	89
Appendix A. dump Option Output	91
ufile_data	91
ufile_name	92
ufile_error	93
ufile_action	93
mdmdb_data	94
mdmdb_error	94
mdmdb_action	95
Glossary	97
Index	107

Figures

Figure 2-1	DMF BFID-Set State Diagram	7
Figure 2-2	Region State Changes for a Partial-State File	8

Tables

Table 2-1	DMF File States	11
Table 2-2	Database Fields	13
Table 2-3	State Change Map for DMF	17
Table 7-1	dmaudit Error Classes	38
Table 7-2	Database Field Descriptions	43
Table 7-3	examine Menu Options	48
Table A-1	ufile_data Output	91
Table A-2	ufile_name Output	92
Table A-3	ufile_error Output	93
Table A-4	ufile_action Output	93
Table A-5	mbmdb_data Output	94
Table A-6	mdmdb_error Output	95
Table A-7	mdmdb_action Output	95

About This Guide

This publication documents DMF recovery and troubleshooting techniques for the Data Migration Facility (DMF) 3.3. For the currently supported release levels, see the `DMF.Readme` file.

Related Publications

See the *DMF Administrator's Guide for SGI InfiniteStorage*.

See the following files on the CD-ROM:

- `/CDROM/platform/DMF.Readme` contains general information about DMF
- `/CDROM/platform/DMF.News` contains a history of features and bug fixes provided with each DMF release
- `/CDROM/platform/DMF.Install` contains installation instructions

For example, for the Solaris platform see:

```
/CDROM/solaris/DMF.Readme
/CDROM/solaris/DMF.News
/CDROM/solaris/DMF.Install
```

See the following man pages:

```
dmattr(1)
dmaudit(8)
dmatvfy(8)
dmfind(1)
dmfsmn(8)
dmfdaemon(8)
dmget(1)
dmhdelete(8)
dmlockmgr(8)
dmmigrate(8)
dmscanfs(8)
dmxfsrestore(8)
dmxfsprune(8)
nohup(1)
```

su(1)
vi(1)
xfsdump(1M) (IRIX)
xfsdump(8) (Linux)
xfsrestore(1M) (IRIX)
xfsrestore(8) (Linux)

Obtaining Publications

You can obtain SGI documentation as follows:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- On IRIX systems, you can use InfoSearch (if installed), an online tool that provides a more limited set of online books, release notes, and man pages. Enter `infosearch` at a command line or select **Help > InfoSearch** from the Toolchest.
- On IRIX systems, you can view release notes by entering either `grelnotes` or `relnotes` at a command line.
- On Linux systems, you can view release notes on your system by accessing the README file(s) for the product. This is usually located in the `/usr/share/doc/productname` directory, although file locations may vary.
- On IRIX and Linux systems, you can view man pages by typing `man title` at a command line.

Conventions

The following conventions are used throughout this document:

Convention	Meaning
command	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures.
manpage(<i>x</i>)	Man page section identifiers appear in parentheses after man page names.

<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
[]	Brackets enclose optional portions of a command or directive line.
...	Ellipses indicate that a preceding element can be repeated.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this publication, contact SGI. Be sure to include the title and document number of the publication with your comments. (Online, the document number is located in the front matter of the publication. In printed publications, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Use the Feedback option on the Technical Publications Library Web page:
<http://docs.sgi.com>
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:
Technical Publications
SGI
1500 Crittenden Lane, M/S 535
Mountain View, California 94043-1351

SGI values your comments and will respond to them promptly.

Introduction

This manual tells you how to use the `dmaudit` command to detect and report every known type of discrepancy between your filesystems and the DMF daemon database, including the following:

- Migrated files for which there are no database entries.
- Database entries for which there are no migrated files.
- Duplicate bit file identifiers (BFIDs) in either the filesystems or the database. (The BFID is the ID assigned to each file during the migration process; it links a migrated file to its data on alternate media.)

The `dmaudit` command is intended primarily for interactive use. It uses a series of scrolling menus to display information and to solicit option selections. However, after you complete the initial configuration, `dmaudit` can also accept the `snapshot` and `report` operations from the command line. This allows `dmaudit` to be run as a background process or in batch mode.

Some of the advantages of using `dmaudit` are as follows:

- `dmaudit` executes while DMF is active. Users can continue to access their migrated files while you simultaneously search for and correct errors.
- `dmaudit` is accurate. `dmaudit` automatically adjusts for any user activity occurring in the filesystems it searches. It therefore is not confused into reporting false errors, nor does it miss errors.
- Discrepancies can be examined interactively. `dmaudit` performs most of its analysis in batch mode and saves the output in indexed files. You can then interactively examine any discrepancy quickly and easily.
- You can decide which errors you want to fix, and you have some control over how a discrepancy is to be fixed.
- You do not need to immediately fix discrepancies. Because `dmaudit` saves all information it needs to fix each error, the errors can be fixed hours or days after they are detected.

Note: The `dmaudit` command is capable of showing you **what** discrepancies exist, but is not able to tell you **why** they happened. Discovering why an error occurred requires some detective work and a considerable knowledge of the internal workings of DMF.

If you do want to determine why a discrepancy occurred, there is information available in some of the `dmaudit` menus to help you narrow down exactly when an error occurred. Examination of daemon logs and journal files may fill in the remaining blanks. Some tips are given in later sections of this manual to help you determine why certain discrepancies occurred.

The Need for `dmaudit`

During normal system operation, the daemon database and the filesystems stay synchronized with each other. Each migrated user file has a unique BFID. Each BFID has one or more active daemon database entries.

However, if a user removes or modifies a file, then the association between that file and the matching database entry is broken because what is in the filesystem no longer matches what is in the database. At this point, the database entry is *soft-deleted*. When a file is soft-deleted, the database entry is marked as inactive and a timestamp is added indicating when the user modified or removed the file. The stale data remains on the alternate media until it is *hard-deleted*, in which the database entry is actually removed.

However, things can get out of synchronization. Examples of some of the discrepancies that might occur are as follows:

- Migrated files that have BFIDs for which no database entries exist
- Active database entries for which no migrated files exist
- Multiple user files that have the same BFID

System crashes are a major source of discrepancies because I/O operations in progress at the time of the crash are not guaranteed to complete successfully. For example, a migrating file might receive a new BFID, but the rewrite of its inode to disk might not succeed. Or perhaps the inode update does complete, but the corresponding database entries are not successfully made.

Inconsistencies also arise if users are allowed to modify or remove migrated files during periods when the daemon is not running, because the kernel is then unable to tell the daemon to soft-delete the corresponding database entries. The unused, or orphan, database entries then accumulate in the database, wasting space on the alternate media.

Use of the `xfsdump` and `xfsrestore` commands can also create inconsistencies. For example, files that have been removed or modified can be restored to their previous state. If the same migrated file is restored multiple times, there will be more than one inode containing the same BFID.

Sometimes the inconsistencies are harmless, or only result in wasted space on the alternate media. In other cases, the discrepancies can prevent a user from accessing one or more files, or can result in the loss of files. `dmaudit` allows the administrator to quickly detect and correct such inconsistencies when they occur, possibly before any data loss becomes permanent.

When to Use `dmaudit`

After `dmaudit` has been initially configured, most sites use `dmaudit` in batch mode on a periodic basis, perhaps once a week. The easiest way to do this is through the use of the `run_audit.sh` script. For more information, see the *DMF Administrator's Guide for SGI InfiniteStorage*.

Sites may also want to use `dmaudit` after a known failure such as a major system crash. Instructions on how to generate a report both interactively and through a `cron` script are provided in Chapter 6, "Detecting Discrepancies" on page 29.

When errors are detected, you should interactively examine and correct them.

Running `dmaudit` on a periodic basis will help you determine why discrepancies have appeared. For example, if your system crashes and a subsequent `dmaudit` run shows discrepancies that were not previously present, you can be reasonably sure that they occurred because of events at the time of the crash.

Disk Cache Manager

The disk cache manager (DCM) for *n*-tier capability presents a class of DMF errors and inconsistencies that `dmaudit` is not capable of detecting or correcting. Until such time that `dmaudit` is enhanced to perform DCM validation, whatever mechanisms

that have been provided to validate the integrity of the DCM will include the guidelines the administrator would normally find in this manual.

Daemon Internals

To use `dmaudit` effectively, you must understand how DMF keeps track of copies of a file's data stored on alternate media, how that data is restored to disk when the user accesses the file, and what happens to those copies when the file is modified or removed.

This chapter discusses the following:

- "BFIDs and BFID Sets"
- "Migration Life-Cycle" on page 6
- "DMF Inode-Resident Data" on page 10
- "DMF Daemon Database Contents" on page 12
- "Errors During Migration" on page 15
- "Remigrating a File" on page 16
- "Removing or Modifying a File" on page 16
- "DMF BFID State Summary" on page 17

BFIDs and BFID Sets

The bit file identifier (BFID) is an object that links a migrated file to copies of its data on alternate media (such as tape). The daemon assigns a unique BFID to each file that it migrates.

A BFID consists of an opaque 16-byte value. (*Opaque* in this context means that the content and format have no fixed definition. The value can be interpreted differently by different processes.)

The BFID represents a unique ID that the daemon inserts into a migrated file's inode and into database entries that point to copies of the file's data. No two migrated files on the same machine should have the same BFID.

For `dmaudit` purposes, a migrated file is one whose inode contains a BFID. A file that does not have a BFID is a nonmigrated file (often referred to in DMF documentation as a *regular file*).

A *BFID set* is the collection of all database entries, special files, and migrated files associated with a particular BFID. The different components that make up a BFID set define its state. During the normal course of migrating and moving files, there are four states that a BFID set can be moved through. Any time that a BFID set contains a combination that differs from one of the four states, it is considered an error.

To understand these BFID-set states, you must understand how DMF uses each of the possible components of a BFID set. "Migration Life-Cycle" on page 6 describes possible states and how files move from one state to the next. The remaining sections describe how the DMF daemon handles migration errors, remigration, and modification of files. "DMF BFID State Summary" on page 17, summarizes the information.

Migration Life-Cycle

For you to be able to understand all the errors and actions reported by `dmaudit`, you must understand the valid bit file identifier (BFID) set states. Figure 2-1 shows a sample of the states that a BFID set can be in during its lifetime; the arrows connecting various states show what actions cause a BFID set to progress from one state to another. Figure 2-2 on page 8 shows an example in more detail of a file in partial-state and the changes to the states of the regions within the file.

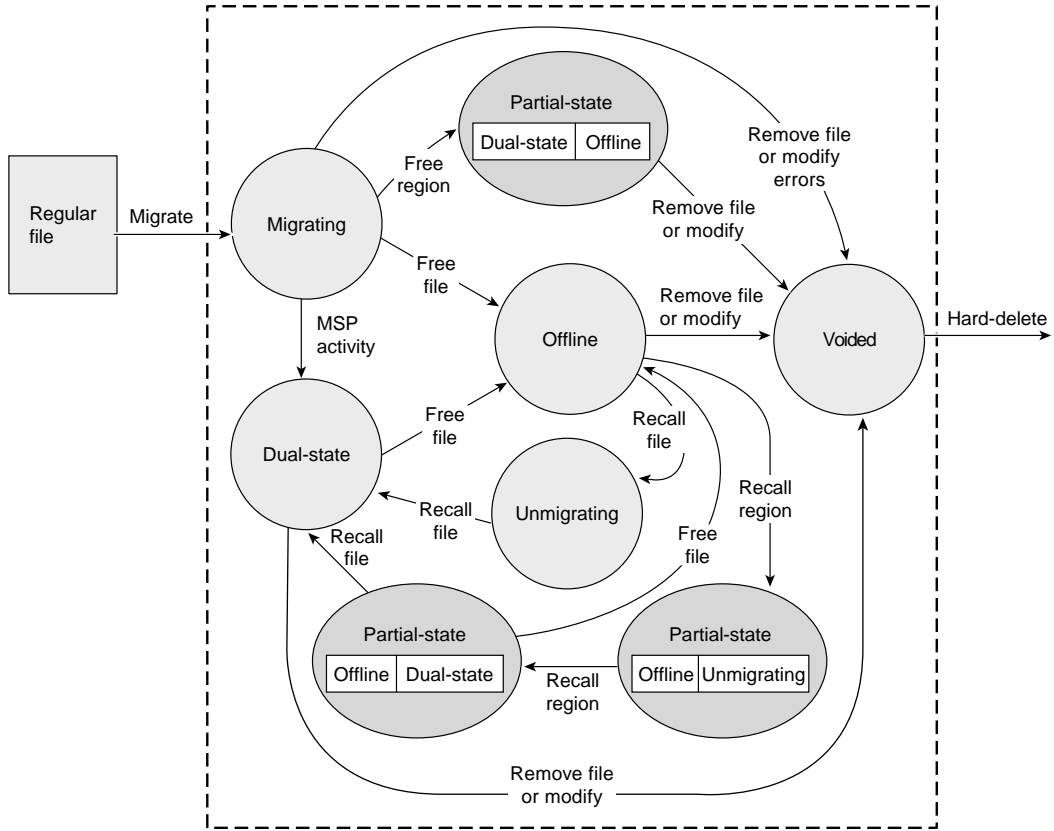


Figure 2-1 DMF BFID-Set State Diagram




Figure 2-2 Region State Changes for a Partial-State File

Recognizing invalid states is the essence of what `dmaudit` does. It collects all information available for a BFID set and attempts to determine which BFID set state best matches the available data. When this has been determined, `dmaudit` then reports as inconsistencies any deviations from the ideal BFID-set state and determines what actions are necessary to return the BFID set to a correct state.

The following sections describe file states and state transitions.

Regular Files

A regular file contains no DMF state information. The file’s inode contains no BFID and there are no MSP or volume group `dmfdaemon` database entries corresponding to the file.

Migrating Files

When a migration request is issued on a regular file (either through a `dmpout` command or through automated space management) the DMF daemon sets the file state to `MIGRATING`, assigns a BFID to the file, and creates an incomplete MSP or volume group `dmfdaemon` database entry containing the BFID and other information relating to the file.

Dual-State Files

When all copies of the file's datablocks have been written to the associated media of their respective MSPs or volume groups, the file's state is set to `DUALSTATE`. This indicates that the file is online and has one or more complete backup copies. In this state, the file's inode contains a BFID. In addition, each copy of the file has a corresponding complete entry in the `dmfdaemon` database.

Similarly, when an offline file is recalled (either explicitly by using a `dmget` command or automatically by a normal filesystem access to the file data such as a `read` system call), the file blocks are retrieved from the backup media and placed back into the filesystem and the file's state is changed to `DUALSTATE`.

Offline Files

If the file's data blocks were released from the filesystem (either due to an automated space management policy or because the owner executed a `dmpout -r` command), the file's state is set to `OFFLINE`. The file's inode contains a valid BFID, the file's underlying data blocks are released, and a complete MSP or volume group `dmfdaemon` database entry exists for the file.

Partial-State Files

A *partial-state file* is a file whose inode contains a BFID and a complete MSP or volume group `dmfdaemon` database entry exists for the file. It contains any mixture of `DUALSTATE`, `OFFLINE`, and `UNMIGRATING` regions.

Voided-State Files

A file backup copy can be voided when any of the following occurs:

- An offline file is modified
- A dual-state file is modified
- A partial-state file is modified
- A migrating file is modified or recalled

In these situations, the BFID is cleared from the file and the `dmfdaemon` database entry is soft-deleted.

DMF Inode-Resident Data

Data migration uses fields within a file's inode to store the BFID, the state of the file, and other information. `dmaudit` retrieves and reports this information when discrepancies are discovered.

The Bit File Identifier (BFID)

The BFID is assigned by the daemon. When a file's inode contains a BFID, it indicates that the file is under DMF control. This field is used as a key into the daemon database to check for validity of the file against the database.

DMF File State Field

The state field shows the current migration state of a file. Table 2-1 lists supported states.

Table 2-1 DMF File States

State	Description
REGULAR	A <i>nonmigrated file</i> ; that is, a file that contains data and whose inode has not yet been assigned a BFID.
MIGRATING	A <i>migrating file</i> ; that is, a file that is in the process of migrating. The inode may or may not contain a BFID.
DUALSTATE	A <i>dual-state file</i> ; that is, a file that contains data and whose inode also contains a BFID pointing to matching copies of that data on alternate media.
OFFLINE	An <i>offline file</i> ; that is, a file whose inode has been assigned a BFID that points to copies of the file's data on alternate media, and all data in the file itself has been removed.
PARTIALSTATE	A file that contains a mixture of OFFLINE, DUALSTATE, or UNMIGRATING regions whose inode also contains a BFID pointing to matching copies of that data on alternate media.
UNMIGRATING	An <i>unmigrating file</i> ; that is, the state in which a file is in the process of being recalled. It holds the data pointers until the MSP or volume group successfully copies the file's data back onto the filesystem disk from alternate media. When the copy is complete, the file is a DUALSTATE file.
NOMIGR	This state is never assigned to user files. The daemon uses it exclusively for unmigration files. The NOMIGR state indicates to programs and to the daemon that this file should never be chosen as a candidate for migration.

Although the daemon is usually responsible for changing the state of a file, an `xfrestore` command may also have to do so when restoring a filesystem from tape. The kernel changes a file's state when the file is modified or removed.

DMF Daemon Database Contents

Four files make up the DMF daemon database. They reside in the daemon's home directory (*HOME_DIR/daemon_name*; *HOME_DIR* is specified in the DMF configuration file; *daemon_name* is the name of the daemon object in the configuration file). The files are as follows:

<code>dbrec.dat</code>	Data file containing the base record information about each file copy.
<code>dbrec.keys</code>	Index file used to access individual records in the data file. The daemon uses the BFID as the index file key; it has the following format: <pre>typedef struct { u_char bfid[16]; } bf_id_t;</pre>
<code>pathseg.dat</code>	Data file containing the overflow path segment blocks for any base (<code>dbrec</code>) records that require extra path field characters.
<code>pathseg.keys</code>	Index file used to access the path segment overflow records for any base (<code>dbrec</code>) records that apply.

The daemon inserts a new entry into its database each time a copy of a file is created. Each database entry contains the same BFID that was stored in the file being migrated.

There are three types of entries in the daemon database, each of which serves a different purpose:

- *An incomplete database entry.* This entry is created when a user file is currently under migration. It represents a database entry with an empty `path` field in the `dbrec` structure (shown following this list).
- *A complete database entry.* This entry is created when the MSP or volume group has completed the file's migration and has entered the file's MSP or volume group key into the `path` field.
- *A soft-deleted database entry.* This entry is created when the `delflag` field of the entry is nonzero, indicating the time that the deletion occurred. Although the entry physically exists in the database, no corresponding user file should exist for the entry. When soft-deleted entries are physically deleted from the database they are said to be *hard-deleted*.

The following `dbrec` structure is used by the daemon to store the contents of one database entry in the `dbrec.dat` file:

```
struct dbrec {
    bf_id_t bfid;           /* off line file bfid */
    dev_t  origdv;         /* original device no. */
    ino_t  origino;        /* original inode number */
    off64_t origsz;        /* original file size (in blocks) */
    time_t otime;          /* original entry time */
    time_t utime;          /* last update time */
    time_t ctime;          /* last check time */
    time_t delflag;        /* delete time */
    int  userid;           /* user id, when archived */
    int  pathlen;          /* length of complete path */
    char  ofilenm[15];     /* original file name (last element) */
    char  proc[8];         /* process name */
    char  path[34];        /* MSP key string */
};
```

Each of the fields in a database entry is explained in Table 2-2. The fields that `dmaudit` uses in its analysis are marked with an X; other fields are not used by `dmaudit` in its analysis because they are not always the most current information and their data does not help `dmaudit` solve inconsistencies. However, you may be able to use them as hints to determine why an error occurred. For example, although it is permissible for a file's ownership to change at any time, in practice that seldom happens, so the user ID of the person who owned the file when it originally migrated is often still the most current information.

Table 2-2 Database Fields

Field	Used by <code>dmaudit</code>	Description
<code>bfid</code>	X	Key by which the daemon accesses its database entries.
<code>origdv</code>		Specifies the device number of the filesystem in which the file resided at the time it was migrated. If you change your disk configuration, the device numbers of your filesystems may change, causing <code>origdv</code> to become out-of-date.

Field	Used by dmaudit	Description
origino		Specifies the inode number of the migrated file at the time it was migrated. This can also become out-of-date if you dump and restore the filesystem that contains the migrated file.
origsz	X	Specifies the size of the migrated file in bytes. Because the size of a file cannot change while it is migrated, this field is always current.
otime		Specifies the date and time when this database entry was first created. This is sometimes called the <i>entry origination time</i> . Usually it corresponds to the time the file first migrated.
utime		Specifies the date and time when any field in this database entry was last changed.
ctime		Specifies the origination time when the database entry was created. Although this can be reset by the administrator using the <code>dmdadm</code> command, in normal practice it never changes.
delflag	X	Specifies whether a database entry is valid or invalid (soft-deleted). This value is initialized to 0 when the database entry is created, and it remains set to 0 as long as the database entry is still valid. When the database entry becomes invalid (for example, if the migrated file was removed), it is soft-deleted by storing the current date and time into <code>delflag</code> in time-stamp format. The database entry cannot be removed immediately because the migrated file might later be restored. You hard-delete the soft-deleted database entries by using the <code>dmhdelete</code> command only when there is no further chance of the file being restored.
userid		Specifies the user ID (UID) of the owner of the file at the time the file was migrated. This can become out-of-date if someone uses the <code>chown</code> command to change the ownership of the file after it has migrated.
pathlen		Specifies the length of the <code>path</code> field described below. This field is used to regulate the creation of <code>pathseg</code> records to contain the overflow characters of the path that will not fit into the <code>path</code> field of this record.
ofilenm		Specifies a null-terminated string composed of up to the first 14 characters of the base name of the file if the pathname was part of the original migrate request. If the pathname was not known at the time the file was migrated, this field contains the string <code>/NONAME</code> .

Field	Used by dmaudit	Description
proc	X	Specifies the name of an MSP or volume group expressed as a null-terminated character string. This field tells the daemon which MSP or volume group to contact to retrieve a copy of the migrated file.
path	X	Specifies either a pathname or a key. For MSP and volume group dmfd daemon database entries, it is the key that the MSP or volume group uses to retrieve a copy of the file. For the volume group, the path field is a key into an LS CAT (catalog) database that contains information about the tape that should be mounted to retrieve the copy. It is a null-terminated string.

The path field of the dbrec base record is a fixed-length field, making the dbrec record a fixed-length record. The path field is supplied by the MSP/LS and can be of any length. The daemon determines the fixed length of the path field of the dbrec record. If the value that the daemon must store in the path field to accommodate the MSP or volume group is longer than the fixed length, the daemon will allocate a sufficient number of pathseg records to hold the overflow path characters.

The pathseg records are keyed from the dbrec bfid and proc fields that uniquely define a dbrec record. When a daemon database record is accessed, any pathseg record path segment extensions are concatenated with the dbrec path field to accurately reconstruct the path value that the MSP or volume group originally supplied to the daemon.

The path field length of the dbrec structure as it is supplied to DMF has a value of 34, which will accommodate the largest dmatls BFID without requiring any overflow pathseg records. If you are running only dmatls, you should not need to adjust the path field of the dbrec structure.

The procedure for adjusting the path field of the dbrec structure is described in *DMF Administrator's Guide for SGI InfiniteStorage*.

Errors During Migration

Sometimes an MSP or volume group may encounter problems that prevent it from completing its copy of a file. For example, the LS could encounter tape errors that prevent it from proceeding. The failing LS then returns an error code to show that it was unsuccessful.

When the daemon receives the error reply, it immediately soft-deletes the MSP's or volume group's `dmfdaemon` database entry (or entries, if more than one MSP or volume group has an entry). The daemon soft-deletes entries by placing the current date and time into a `delflag` field.

The daemon must then *void the BFID* (remove the BFID from the user file inode). This is done because the user file could not be migrated exactly as requested.

The daemon then sets the file's state to `REGULAR`.

Remigrating a File

The process of migrating a dual-state file is sometimes known as *remigrating* the file. There is only one minor difference between remigrating a file and migrating a file for the first time: the amount of time the migration takes.

When a dual-state file is remigrated, the daemon checks to see whether a copy of the file already exists on each of the requested MSPs and volume groups. If so, the daemon does not have to instruct the MSP or volume group to make a new copy of the file, because the old copy is still valid.

Unless the DMF configuration has changed, a dual-state file always migrates to the same MSPs and volume groups to which it originally migrated. This means that the BFID set only spends a few milliseconds in the incompletely migrated state before advancing to the fully migrated state.

If the dual-state file is being migrated to a new MSP or volume group, a considerable amount of time is spent in the incompletely migrated state while waiting for the new MSP or volume group to make a copy of the file.

Removing or Modifying a File

When a user either modifies or removes a file that contains a BFID, the kernel notifies the daemon by generating an event for the file that the daemon has registered to receive. If the user modified the file, a write event is generated and the daemon will remove the BFID from the inode and change the file's state to `REGULAR`, because the copies of the file are no longer current. If the user removed the file, a destroy event is generated and the daemon will not need to do anything to the file (because it is no longer present).

In either case, the daemon will soft-delete all remaining database entries for that BFID to indicate that those copies are no longer valid.

DMF BFID State Summary

Table 2-3 shows a summary of the information presented in previous sections.

Table 2-3 State Change Map for DMF

DMF BFID-Set State	User File State	MSP or Volume Group Entries	Entries Soft-Deleted?
Incompletely migrated	Migrating	At least one incomplete ¹	No
Fully migrated	Dual-state	All complete	No
Freed	Offline	All complete	No
Incompletely unmigrated	Unmigrating	All complete	No
Mixture of fully migrated, freed, and incompletely unmigrated	Partial-state	All complete	No
Voided	2	Either complete or incomplete	Yes

`dmaudit` uses the information in this table when it examines all the information available for each BFID set. Any BFID set that exactly matches one of the entries in the appropriate table is considered error-free.

If `dmaudit` finds a BFID set that does not fit exactly into one of the entries in the table, it is reported as having errors and `dmaudit` lists the actions necessary to remove the inconsistencies.

¹ In the incompletely migrated state, complete MSP or volume group entries may also exist.

² The file state is regular or the file has been removed.

How `dmaudit` Detects Errors

The `dmaudit` command collects and analyzes information from the filesystems and the daemon database. It then stores that information in a directory known as the *working directory*. This process is known as *taking a snapshot* and is normally done in batch mode. `dmaudit` takes a snapshot by performing the following actions:

1. Scans filesystems for migrated files. Each filesystem that might contain migrated files is searched for files that contain a BFID. Information about the state of each of these files is saved in the working directory. At the same time, `dmaudit` instructs the daemon to maintain a log of all changes that occur to migrated files while the search is in progress.
2. Copies the daemon database to the working directory. When the filesystem scans are complete, `dmaudit` asks the daemon to place a copy of the daemon database into the working directory together with the log of changes that occurred during the filesystem scan.
3. Brings the filesystem information up-to-date. The data collected during the filesystem scan is updated using the migrated file change log supplied by the daemon. When completed, this makes the filesystem information appear as if it were collected at the same instant that the daemon database copy was made. The filesystem and database data files are now synchronized with each other.
4. Sorts and merges the filesystem and database information. The filesystem information and database entries are each sorted into BFID order and are then merged together into a large file stored in the working directory.
5. Analyzes each BFID set for errors. All data available for each BFID set in the composite file is examined for inconsistencies. `dmaudit` compares the state in the user file's inode against any daemon database entries and flags any inconsistencies. It also computes and saves information on how to correct each error.
6. Issues a report. The report gives a brief summary of the kinds of errors encountered.

`dmaudit` retains in its working directory all snapshot information for each run. After the report has been issued, you have the option of using `dmaudit` interactively to examine each BFID set that contains errors.

For examples, see "The snapshot Option" on page 30 and "Example of a Report with Discrepancies" on page 34.

Using `dmaudit` Interactively

The `dmaudit` command is intended primarily for interactive use; some steps, such as configuration, must be done interactively. `dmaudit` uses a series of menus to display information and to solicit input from you.

Environment Variables

For `dmaudit` to display its menus properly, it needs to know the dimensions of your screen. By default, it uses the values found in the environment variables `LINES` and `COLUMNS`. If the variables are not defined, `dmaudit` assumes a screen length of 24 lines and a width of 80 columns. You can either set these variables manually or use the `resize` command; see the `resize(1)` man page for further information. C shell users normally enter something similar to the following to set the screen size environment variables:

```
eval `resize -c`
```

Sometimes `dmaudit` invokes an editor to request additional information from you. It uses the editor defined in the environment variable `VISUAL` (which is read first) or `EDITOR`; if neither is defined, it uses the `vi` command by default. If you want to use an editor other than the default, set one of the two environment variables before starting `dmaudit`. For example, a C shell user who wants to use the `emacs` editor could enter the following:

```
setenv VISUAL emacs
```

Page-Wait Mechanism

Occasionally `dmaudit` needs to display more information than can fit on your screen. In these cases, it uses a page-wait mechanism to retain the current screen of data until you are ready to continue. Your screen fills with data down to the last line, which appears as follows:

```
Enter <CR> to continue:
```

When you are done examining the current screen, press `ENTER` to view the next screen of data. If you are not interested in the remaining data, you may enter the first

character of the next menu item and press `ENTER` to skip directly to processing the next menu item. There is no way to page back to a previous screen after you have pressed `ENTER`.

The page-wait feature is needed for users with terminals whose data would otherwise scroll off the screen before they could read it. If you are using a windowing system such as the X Window System environment, you can minimize the occurrence of page waits by running `dmaudit` in a window that contains as many lines as possible and by using the scroll bars available in the X Window System environment.

Setting the Initial Configuration

The first time that you run `dmaudit`, you must configure it; this step must be done interactively. This consists of the following:

- Giving `dmaudit` a list of all the filesystems that might contain migrated files
- Specifying the location of a working directory where `dmaudit` can keep snapshot data for indefinite periods of time.

The `dmaudit` command saves the configuration information that you enter in the following file, where `HOME_DIR` is the DMF home directory you specified in the DMF configuration file:

```
HOME_DIR/dmaudit_dir/checkpoint
```

The information is used in all subsequent invocations of the command.

This chapter discusses the following:

- "Providing a Working Directory"
- "Specifying Filesystems to Scan" on page 24

Providing a Working Directory

To configure `dmaudit`, you must first either log in as `root` or use the `su` command to become the super user. To execute the command interactively, enter the following:

```
dmaudit
```

The following display appears the first time that you invoke `dmaudit`:

```
This program must create a working directory named 'working_dir' to hold several very large work files. The files may be needed in subsequent executions of this program, so the directory should be placed in a file system that is not cleared frequently.
```

```
Please enter the full path name of an existing directory in which subdirectory 'working_dir' can be created (<CR> to quit):
```

Enter the name of a directory, such as `/tmp`. If the directory you specified does not exist or is otherwise invalid, `dmaudit` describes the nature of the problem and prompts for a new pathname.

If you decide not to continue, you can exit the program at this time by pressing `ENTER`.

Specifying Filesystems to Scan

After you have entered the directory name, `dmaudit` asks for the list of filesystems it should search when looking for migrated files. It presents the following menu:

You must select the file systems to be scanned when this program searches for migrated files. By default the list contains all file systems currently mounted.

IF THERE ARE ANY FILE SYSTEMS CONTAINING MIGRATED FILES THAT DO NOT APPEAR IN THIS LIST, EXIT THIS PROGRAM AND MOUNT THEM BEFORE CONTINUING!!!

Failure to do so could mean the loss of the data in those files, because this program sometimes removes database entries if it cannot find matching migrated files for them. File systems that you are sure do not contain migrated files may be removed from this list to speed execution. Remember that using `restore` to load files dumped from a DMF-configured file system can result in migrated files in other file systems. If there is any question whether a file system contains migrated files, leave it in the list.

Select:

```
<view>    View the current file system scan list
<edit>    Edit a fresh copy of the file system scan list
<accept>  Accept the current file system scan list
<quit>    Quit
```

Please enter your selection:

Enter `view` to display the list of filesystems that are currently mounted and that support data migration. (To support data migration, a filesystem must be mounted with the `dmi` option on IRIX systems and with the `dmapi` and `mtpt = mountpoint` options on Linux systems. For information about mounting filesystems, see the *DMF Administrator's Guide for SGI InfiniteStorage*).

When the list has been shown, `dmaudit` automatically redisplay the filesystem selection menu. You should see something like the following:

```
/                /core
/admin           /mnt_tmp
/cloudy/ccn      /tmp
/cloudy/mktg     /usr
/cloudy/sdiv/comp /usr/adm
/cloudy/sdiv/lib /usr/dm
/cloudy/sdiv/net /usr/spool
/cloudy/sdiv/qte /usr/src
/cloudy/sdiv/xfiles /usr/tmp
```

Select:

```
<view>      View the current file system scan list
<edit>      Edit a fresh copy of the file system scan list
<accept>    Accept the current file system scan list
<quit>      Quit
```

Please enter your selection:

If you have many filesystems mounted, they may not all fit on the screen at the same time, in which case the last line on the screen instead looks like the following:

Enter <CR> to continue:

If that is the case, then when you are done examining the current screen, press `ENTER` to view the remainder of your filesystems.

Normally you should choose the `accept` selection to run with all `dmi`-mounted filesystems included in the audit. This is the proper method to use when you are unsure which filesystems contain migrated files. `dmaudit` saves the pathname, device name, and filesystem type of each filesystem in the list in its checkpoint file and uses this list in all future snapshots.

You might want to remove some of the filesystems from the list if you know they will never contain migrated files because of the following reasons:

- Filesystems take time to scan. If you can remove filesystems from the list, you will reduce `dmaudit` execution time.
- Some filesystems change frequently. `dmaudit` expects to find each filesystem in the list mounted in the same location every time it is run, and does not allow you

to take a snapshot or fix discrepancies if any filesystem has changed. You may want to remove filesystems that are not always mounted.

The important thing is to make sure that any filesystem that might contain migrated files is in the list. `dmaudit` does not allow you to add filesystems to the list that are not currently `dmi-mounted`, so if any such filesystems do not appear in the list, enter `quit` to exit the program and mount them before continuing.

To continue the example, assume the following:

- You do not want to scan any of the filesystems shown in the previous screen starting with `/usr` because you know that they are large filesystems and that no migrated files will ever reside in them
- No files will ever be migrated in `/`
- The `/mnt_tmp` filesystem is only used as a place to mount temporary filesystems
- The `/tmp` filesystem is only used for temporary files that are not migrated
- The `/` (root) filesystem should be removed from the list. (Sites that frequently switch root filesystems when building new systems may want to remove the root filesystem from their list.)

If you want to exclude these filesystems, enter `edit`. This puts your screen under control of the editor defined in the `VISUAL` or `EDITOR` environment variable, (`VISUAL` takes precedence); if both are undefined, it uses `vi`. You are shown a list of filesystems in alphabetical order, one per line.

For the previous example, you would see the following:

```
/
/admin
/cloudy/ccn
/cloudy/mktg
/cloudy/sdiv/comp
/cloudy/sdiv/lib
/cloudy/sdiv/net
/cloudy/sdiv/qte
/cloudy/sdiv/xfiles
/core
/mnt_tmp
/tmp
/usr
```

```
/usr/adm
/usr/dm
/usr/spool
/usr/src
/usr/tmp
~
~
```

Using the editor, remove the lines associated with filesystems that you **do not** want to include in the audit. Next, save your changes and exit the editor.

After you exit the editor, `dmaudit` reports any errors that it found in your input and then redisplay the filesystem selection menu. If errors are detected, `dmaudit` discards your entire edit session, and you must select `edit` again and reenter all your changes.

It is a good idea to select the `view` option one last time to verify that `dmaudit` accepted your selections.

In the current example, you would see the following:

```
/admin                /cloudy/sdiv/net
/cloudy/ccn           /cloudy/sdiv/qte
/cloudy/mktg         /cloudy/sdiv/xfiles
/cloudy/sdiv/comp    /core
/cloudy/sdiv/lib
```

Select:

```
<view>      View the current file system scan list
<edit>      Edit a fresh copy of the file system scan list
<accept>    Accept the current file system scan list
<quit>      Quit
```

Please enter your selection:

The unwanted filesystems no longer appear in the list.

Once you are satisfied with your filesystem selections, enter `accept` to continue. `dmaudit` responds with the following menu:

MAIN MENU

Select:

<code><snapshot></code>	Take a snapshot and report status of file systems and databases
<code><config></code>	Examine or modify configuration information
<code><quit></code>	Quit

Please enter your selection:

Configuration of `dmaudit` is now complete. At this point you may quit, enter `config` to examine certain other optional configuration parameters (see Chapter 8, "Changing the Configuration" on page 83), or enter `snapshot` to initiate a snapshot (as described in "The snapshot Option" on page 30).

Detecting Discrepancies

The next step in using `dmaudit` is to generate a report that summarizes discrepancies between the filesystems and the DMF daemon database. As was mentioned earlier, this is also known as *taking a snapshot*. The DMF daemon `dmfdaemon` must be running before you can attempt this step.

This chapter shows you how to generate the report both interactively and in batch mode:

- "Snapshot Resource Requirements"
- "Taking a Snapshot Interactively" on page 30
- "Example of a Report with Discrepancies" on page 34
- "Taking a Snapshot in Batch Mode" on page 35

You should read the entire chapter before generating a report on your own machine so that you can choose the method that best fits your needs.

Snapshot Resource Requirements

Taking a snapshot with `dmaudit` can require a large amount of both wall-clock time and CPU resources, depending on the size and number of files being scanned. During the snapshot, `dmaudit` must sort and merge several very large data files.

The first time you take a snapshot, you should do it during a time period in which you know `dmaudit` is able to execute continuously for several hours. Keep track of how long it takes to generate the report so that you can determine how much time to allot in future runs. If the program is stopped for any reason before the report generation is complete, the entire snapshot is discarded and you must start over.

Taking a Snapshot Interactively

This section describes the various options available to you when you take a `dmaudit` snapshot interactively.

If you just finished the initial configuration of `dmaudit`, you should already be positioned at the `MAIN` menu. If you are not currently running `dmaudit`, enter the `dmaudit` command.

If you are running `dmaudit` but are positioned at some other menu, enter `up` as many times as necessary until you arrive at the `MAIN` menu. Your screen should look like the following:

```
MAIN MENU
```

```
-----
```

```
Select:
```

```
<snapshot>  Take a snapshot and report status of file systems and databases
<config>    Examine or modify configuration information
<quit>      Quit
```

```
Please enter your selection:
```

The snapshot Option

To generate a report of all discrepancies in your filesystems and daemon database, enter `snapshot` from the `MAIN` menu. A display similar to the one shown in the following example eventually appears:

```
DAEMON DATABASE ERROR REPORT
```

```
-----
```

```
No errors were discovered comparing the file systems against the daemon
database.
```

```
DATA MIGRATION CONFIGURATION
```

```
-----
```

```
Data migration home directory:      /dmf/home
```

```
Data migration binaries directory:   /usr/sbin
```

```
Server name:                         daemon
```



```
Server home directory:      /dmf/home/daemon
Server spool directory:    /dmf/spool/daemon
```

```
Data migration daemon process ID:      1258240
```

```
MSP name   MSP type
-----   -
dcm_msp    dmdukmosp
```

Enter <CR> to continue: **ENTER**

```
Name      Type
-----   -
ls1       dmatls
vg9a20    volume group   in Library Server ls1
vg9a01    volume group   in Library Server ls1
vg9a02    volume group   in Library Server ls1
```

LIST OF FILE SYSTEMS SCANNED

```
-----
/dmf1     xfs
/dmf2     xfs
/dmf3     xfs
/dmf4     xfs
```

Enter <CR> to continue:

The report shows the following:

- The database error report
- The current DMF daemon, MSP, and/or volume group configuration
- The list of filesystems that were searched for migrated files. The type of each filesystem (such as xfs) is shown next to its pathname.

Note: If the report is longer than one screen (as is shown in this example), press ENTER to view the remainder of the report.

In this example, no errors were detected in either the filesystems or the daemon database.

When you press ENTER at the end of the report, `dmaudit` returns to the MAIN menu, which then appears as follows:

MAIN MENU

Select:

<inspect>	Inspect and correct file system and database errors
<report>	Reprint status report for the current snapshot
<verifymsp>	Check the dmatls tape library server databases against the daemon databases
<snapshot>	Take a snapshot and report status of file systems and databases
<free>	Release all file space used by the current snapshot
<config>	Examine or modify configuration information
<quit>	Quit

The `inspect`, `report`, `verifymsp`, and `free` options are described in the following sections.

The inspect Option

The `inspect` option is normally used only in cases in which discrepancies have been detected. It allows you to examine each error in detail and to correct those errors. Even if there are no discrepancies, `inspect` can be used to examine individual BFID sets in the daemon database.

The report Option

The `report` option allows you to reprint the report from your most recent snapshot. Because all information contained in the report is determined during the snapshot phase, the `report` option can reissue the report instantaneously. The report looks the same as the one shown in "The snapshot Option" on page 30.

This option is useful if you took a snapshot earlier and want to view its report again. It is also handy if you normally take snapshots in batch mode. If a batch run indicates that there are errors, you can then execute `dmaudit` interactively and read the report directly.

The `verifymsp` Option

The `verifymsp` option lets you tell `dmaudit` to do the following:

- Run the `dmadvfy` command using the data captured during the snapshot
- Adding the information to the `dmaudit` report

A display similar to the following eventually appears:

```
Please enter your selection: verifymsp
(06/01/04 10:28:19) Using dmaudit databases /tmp/working_dir.
(06/01/04 10:28:19) Extracting dmf records.
(06/01/04 10:28:19) Verifying library server ls.

LIBRARY SERVER ls ERROR REPORT
-----

  Bitfile ID 40b4b640000000000000000003 - for vsn SGI007 chunk 1 entry for
    volume group vglA missing from daemon db.
  Bitfile ID 40b4b640000000000000000004 - for vsn SGI007 chunk 2 entry for
    volume group vglA missing from daemon db.
  Bitfile ID 40b4b640000000000000000005 - for vsn SGI007 chunk 4 entry for
    volume group vglA missing from daemon db.
3 Errors found in CAT database.
0 Errors found in VOL database.

MAIN MENU
-----
```

You can view the output by using the `report` option, which is described in "The `report` Option" on page 32.

The `free` Option

After you take a snapshot, a potentially large amount of disk space remains allocated to files in the working directory to contain all the information `dmaudit` needs to generate its report. This is done so that you can examine the snapshot at your leisure. The disk space normally remains reserved until you take your next snapshot.

If no errors were detected in the snapshot, however, you might want to enter the `free` option to release that disk space so that it can be used for other purposes until

the next time you take a snapshot. The `free` option discards all snapshot files while retaining all configuration information.

After the snapshot files have been removed, the MAIN menu once again looks like the following:

```
MAIN MENU
-----
```

Select:

```
<snapshot>  Take a snapshot and report status of file systems and databases
<config>    Examine or modify configuration information
<quit>      Quit
```

Please enter your selection:

You do not have to manually select `free` each time before taking a new snapshot. The `snapshot` option automatically releases all space from a previous snapshot before taking a new one.

Example of a Report with Discrepancies

If discrepancies are detected between the filesystems and the daemon database, `dmaudit` attempts to summarize them in the last section of its report.

The following is an excerpt from a report that detected several kinds of errors:

```
DAEMON DATABASE ERROR REPORT
-----
```

```
There are 14 bitfile IDs in use by more than one file that cannot be
corrected without additional information from you.
```

```
There are 3 bitfile IDs in use by more than one file that can be
automatically corrected.
```

```
There are 2 user files whose data cannot be recovered.
```

```
There are 231 user files that have correctable errors.
```

```
There are 18 bitfile IDs in the daemon database for which no user files
```

can be found.

There are 43 bitfile IDs whose errors cannot be corrected until a new snapshot is taken.

There are 12 bitfile IDs with files that cannot be recalled or migrated.

There are 9 bitfile IDs with files that are internally inconsistent.

`dmaudit` summarizes the errors that it finds into nine general classes (described in Table 7-1 on page 38). `dmaudit` allows you to examine each of the error classes separately.

Taking a Snapshot in Batch Mode

To avoid tying up your terminal for long periods of time, `dmaudit` also allows you to specify the `snapshot` option in batch mode. To take a snapshot in batch mode, specify `snapshot` as a parameter on the command line. For example, a Bourne shell user might enter the following:

```
nohup dmaudit snapshot >rpt 2>err &
```

(The `nohup` command allows `dmaudit` to continue processing even if you log out of the system.) When the snapshot completes, the report is written to standard output, in this case to the file `rpt`. If `dmaudit` encounters fatal errors, it issues messages to standard error (in the example, this is the file `err`).

The exit status from a snapshot run in batch mode can be used in shell scripts to determine whether errors were detected, as follows:

- An exit status of 0 indicates that `dmaudit` completed the snapshot and that no filesystem or database errors were detected.
- An exit status of 1 indicates that filesystem or database errors were discovered by `dmaudit`.
- Any other exit status indicates that `dmaudit` aborted with a fatal error.

The `report` and `free` options may also be used in batch mode. The following is an example cron script that could be used to take snapshots on a periodic basis:

```
#!/bin/sh
#      script to take a snapshot and report errors by mail
#

TMPDIR=/TMP
MAILLIST="root"

dmaudit snapshot 2>${TMPDIR/errmsg 1>/dev/null
STATUS=$?

if [ $STATUS -eq 0 ]
then
    dmaudit free      # remove unneeded files
    exit 0
fi

if [ $STATUS -eq 1 ]
then
    dmaudit report | mailx -s "dmaudit error report" $MAILLIST
else
    cat ${TMPDIR/errmsg | mailx -s "dmaudit failed with this msg" $MAILLIST
fi
exit 1
```

If the snapshot completes successfully with no errors detected, the script uses the `free` option to release all working directory space. If errors are detected or if `dmaudit` fails, mail is instead sent to the administrator.

Examining and Correcting Discrepancies

This chapter provides an overview of how to use `dmaudit` to examine and correct errors and a walk-through for each menu, using as an example some of the inconsistencies introduced when a filesystem that uses DMF is restored. It also includes additional information on each of the error classes. It contains the following sections:

- "Overview of the Correction Process"
- "The `INSPECT` Menu" on page 39
- "Fixing BFID Sets for Which No User Files Exist" on page 45
- "Fixing Files with Correctable Errors" on page 50
- "Cleaning Up Files with Unrecoverable Errors" on page 56
- "Cleaning Up Files with the Same BFID but Different Sizes" on page 63
- "Cleaning Up Multiple Files with the Same BFID and Size" on page 67
- "Cleaning Up Multiple Files with the Same BFID and Size" on page 67
- "BFID Sets That Cannot Be Immediately Corrected" on page 78
- "The `apply` Option" on page 80

Overview of the Correction Process

Correcting errors with `dmaudit` uses a two-step process. The first step is to interactively select which errors you want to correct; the second step is to correct those errors in either interactive or batch mode.

Errors are divided into eight major classes. Each of these classes has its own correction menu, although the contents of these menus are largely the same. The error classes are described in Table 7-1.

Table 7-1 `dmaudit` Error Classes

Class	Description
1	More than one file has the same BFID, and <code>dmaudit</code> cannot determine by itself which database entries go with which migrated file. User data may or may not be lost.
2	More than one file has the same BFID, but <code>dmaudit</code> has enough information to determine which database entries go with which migrated file. User data may or may not be lost. This class implies database errors.
3	Migrated files for which data cannot be found.
4	Files with minor errors that can be repaired without loss of user data. These are usually caused by events such as the restoration of a filesystem. No user data is at risk in a recoverable error, so they are usually corrected with detailed examination.
5	Active daemon database entries for which no migrated files can be found. No user data is affected by this type of error. The only negative effect of errors in this class is the space wasted by the unused file copies on the migration media. This type of error usually occurs when there is user activity while the daemon is not running.
6	BFID sets with errors for files with one of the following conditions: <ul style="list-style-type: none"> • Their file states or daemon database entries were actively changing while the snapshot was in progress • The files have changed since the snapshot completed The status of these files cannot be evaluated and therefore they cannot be fixed without taking another snapshot. You can, however, look at the most recent state of the BFID sets.
7	Unused
8	BFIDs found in filesystems not configured for DMF and for which no database entries were found. User data has been lost.
9	BFID sets with files that are internally inconsistent.
10	Partial-state files for which a complete copy of the data cannot be found.

Selecting Which Errors to Correct

To select errors to be corrected, enter the number of the error class you want to correct, which displays the appropriate error-class menu. Each error class must be corrected separately. Within each error class menu, choose which errors you want to correct by accepting the modifications `dmaudit` needs to correct them. Each error

class except class 6 has an `accept` option for this purpose. If you change your mind, you can cancel your acceptance by using the `cancel` option.

Correcting the Errors

When you are satisfied with the list of error corrections you have accepted, the next step is to return to the `MAIN` menu and apply your changes using the `apply` option. There are two ways to do this:

- If the corrective actions presented by `dmaudit` imply that very little tape and disk activity is required, you may want to fix the errors immediately. In that case, select `apply` on the `MAIN` menu and wait. When the `MAIN` menu next appears, all errors are corrected.
- If the corrective actions imply that a lot of tape and disk activity is required, you may want to wait until off-peak hours to actually correct the errors. You can perform the corrections interactively or you can do it in batch mode by specifying `apply` as a parameter on the `dmaudit` command line.

The `INSPECT` Menu

To introduce the most common correction options available within `dmaudit`, the next few sections show an example in which the `xfsdump` and `xfrestore` commands were used to cause errors to appear. Restoring a filesystem can cause errors because the filesystem is being returned to a previous state, but the daemon database remains in its current state.

An attempt was made in these examples to produce as many different kinds of errors as possible. If you have to restore a filesystem, your next `dmaudit` report will probably only contain a subset of the errors shown here.

Errors were produced first by using a `dump` command to dump a DMF filesystem. User activity was then simulated by migrating, recalling, and removing files in the filesystem. The filesystem was then restored using a `restore` command, and `dmaudit` was run to produce an error report.

For purposes of the example, it is assumed that you have already taken a snapshot and that errors were detected. The error summary portion of the `dmaudit` report would show something like the following:

DAEMON DATABASE ERROR REPORT

There are 8 user files that have correctable errors.

There are 2 bfids in the daemon database for which no user files can be found.

Assume you decide to correct the errors. You must always start from the MAIN menu screen. If you are not currently running `dmaudit`, issue the `dmaudit` command.

If you are running `dmaudit` but are in some other menu, enter up as many times as necessary until you arrive at the MAIN menu.

Your screen will look like the following:

MAIN MENU

Select:

- <inspect> Inspect and correct file system and database errors
- <report> Reprint status report for the current snapshot
- <verifymsp> Check the dmatls tape library server databases against the daemon databases
- <snapshot> Take a snapshot and report status of file systems and databases
- <free> Release all file space used by the current snapshot
- <config> Examine or modify configuration information
- <quit> Quit

To examine and correct errors, enter `inspect`, which causes the following menu to appear:

INSPECT FILE SYSTEM AND DATABASE ERRORS

Errors are divided into several major classes. Only those classes that pertain to this snapshot are displayed.

Select:

- <4> Examine files that have correctable errors

```

<5>      Examine bfid's in the daemon database for which no user files
          can be found
<bfid>   Examine all files and database entries for any bfid you specify
<search> Scan file systems for names of all files with errors (very slow)
<up>     Return to the previous menu

```

Please enter your selection:

There is a direct correspondence between the two error classes shown in this menu and those mentioned in the error report. You must examine and correct each of these classes separately. To examine a particular class, enter its class number at the prompt. See "Fixing BFID Sets for Which No User Files Exist" on page 45.

The following sections describe the `bfid` and `search` options; the `up` option returns you to the MAIN menu. The error-class options (4 and 5) are described in "Fixing BFID Sets for Which No User Files Exist" on page 45, and "Fixing Files with Correctable Errors" on page 50.

The `bfid` Option

The `bfid` option is used to examine everything that is known about a particular BFID set. You can enter a particular BFID set value or select the default value displayed at the prompt:

Select:

```

<bfid>   Examine all files and database entries for any bitfile ID you
          specify
<search> Scan file systems for names of all files with errors (very slow)
<up>     Return to the previous menu

```

Please enter your selection: **bfid**

(34745a2b000000000000000444): <enter an alternate bfid here or return>

BITFILE ID = 34745a2b000000000000000444

ENTRIES IN USE:

```

01. dual-state user file - dev 50331653 (daemon), fhandle
    010000000000001888c5f6086f39b65d000e00000000000000000000401f8c, uid 285,
    size 175538
02. daemon MSP/VG <ftp> database entry - not soft deleted, size 175538, key

```

7: Examining and Correcting Discrepancies

<abc/34745a2b0000000000000444>

Select:

- <next> Examine the next bfid in the snapshot
- <prev> Examine the previous bfid in the snapshot
- <mode> Switch to full display
- <dump> Append all information about this bfid to a text file in a format suitable for machine processing
- <up> Return to the previous menu

Please enter your selection:

This is called an abbreviated display because it shows only some of the information known about this BFID set. It shows only those pieces of information most commonly examined and has the advantage of taking up less screen space. There are two objects associated with BFID 34745a2b0000000000000444: a user file and an MSP or volume group daemon database entry. The user file has a state of dual-state, a device number of 50331653, a user identification (UID) of 285, and a size in bytes of 175538. The database entry is for an MSP or volume group named ftp, and the key the MSP or volume group uses to retrieve the file is abc/34745a2b0000000000000444. Furthermore, the soft-delete field is 0, indicating that the entry is still active.

You can see a more detailed description of the BFID set by selecting the mode option, which redisplay the BFID set in full-display mode. This display shows all information known by `dmaudit` about the BFID set. If you enter mode for this BFID set, you see the following (the mode option is actually a toggle; each time you select it, the display switches to the display type not currently shown):

BITFILE ID = 34745a2b0000000000000444

ENTRIES IN USE:

01. dual-state user file - dev=50331653 (daemon)
fhandle=01000000000001888c5f6086f39b65d000e000000000000000000401f8c
size=175538 uid=285 nlinks=1
02. daemon MSP/VG database entry - dev=50331653 ino=4202380 size=175538 uid=285
otime=Nov 20 13:08:36 1997 utime=Nov 20 13:08:36 1997 ctime=Nov 20 13:08:36
1997 dtime=0 name=dmf_tst.00061 msp/vg=ftp key=<abc/34745a2b0000000000000444>

Select:

```

<next>  Examine the next bfid in the snapshot
<mode>  Switch to abbreviated display
<dump>  Append all information about this bfid to a text file in a format
         suitable for machine processing
<up>    Return to the previous menu

```

Please enter your selection:

Several new pieces of information appear. The number of hard links for the user file are now shown, as is every field in the daemon's MSP or volume group database entry. In both cases, information that does not fit on the current line is continued on subsequent lines without line numbers.

To conserve screen space, the names of the fields in a database entry in this display are somewhat shorter than the names actually used by the daemon in its dbrec structure (see "DMF Daemon Database Contents" on page 12, for a description of the structure). Table 7-2 shows the correspondence between names for the fields in the full display and in the daemon dbrec structure.

Table 7-2 Database Field Descriptions

Display Name	dbrec Structure	Description
dev	origdv	Device number of the file
ino	origino	Inode number of the file
size	origsz	File size in bytes
uid	userid	User ID of the file owner
otime	otime	Origination time of the entry
utime	utime	Last update time of the entry
ctime	ctime	Last check time of the entry
dtime	delflag	Soft-delete time of the entry
name	ofilenm	Base name (if known) of the file
msp	proc	MSP or volume group name
key	path	MSP or volume group key or pathname

Several other options are available on the `bfid` menu. Because `dmaudit` maintains its data in BFID-set order, the `next` and `prev` options allow you to examine BFID sets in the database adjacent to the current one.

Entering `next` shows you the next higher BFID set in use. If you enter `prev`, you would see consecutively lower BFID sets. This can sometimes be useful if you want to see the state of BFID sets whose files were migrated at about the same time as the current file.

The `dump` option is not commonly used. It allows you to dump all information contained in the full display into a file in machine-readable form. If you select the `dump` option, it prompts you for the name of a file:

```
File to which you want the dump appended?
```

Enter the pathname of a file to which you want the information to go. Either full or relative pathnames can be used. If the file does not already exist, `dmaudit` creates it. If it does exist, it appends the information to the end of the file. If you instead press `ENTER`, `dmaudit` assumes that you have changed your mind and returns to the `bfid` menu.

The format of the dump information is complex; for a description, see Appendix A, "dump Option Output" on page 91.

When the dump completes, `dmaudit` redisplay the `bfid` menu.

From this display, enter `up` to return to the `INSPECT` menu.

The search Option

The `search` option only appears on the `INSPECT` menu when there are BFID sets with errors that have user files associated with them. For example, if there were only class-5 errors in the following menu, `search` would not appear because no user files would be involved:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
```

```
-----  
Errors are divided into several major classes. Only those classes that pertain  
to this snapshot are displayed.
```

```
Select:
```

```
<4>      Examine files that have correctable errors  
<5>      Examine bfid's in the daemon database for which no user files
```

```
can be found
<bfid>      Examine all files and database entries for any bfid you specify
<search>    Scan file systems for names of all files with errors (very slow)
<up>       Return to the previous menu
```

Please enter your selection:

The purpose of `search` is to find the pathnames for each user file whose BFID set has errors associated with it. `dmaudit` does this by recursively scanning through the directories of each filesystem that contains BFID sets with errors.

Note: This process can take a long time, perhaps hours on very large filesystems. You can interrupt the process by pressing `Control-C`, but then you must start the search again from the beginning. It is better to examine the errors first to see if you really need to know the names of the files. If you later decide you need the names, you can return here and enter `search`.

A search from the `INSPECT` menu finds the names of all user files whose BFID sets have errors for all applicable classes. The individual error class menus also have a `search` option, but their search is restricted to those user files that are in their error class. If you need names for multiple classes, it is usually more efficient to get all the names at once from the `INSPECT` menu. If you enter `search`, `dmaudit` pauses for an indefinite period of time as the filesystems are scanned. When the `INSPECT` menu next reappears, all the names have been collected, and the `search` option disappears from the menu.

Fixing BFID Sets for Which No User Files Exist

To continue with the example, you are now ready to examine each of the error classes to see what kinds of errors occurred. Assume that you entered 5 to look at those BFID sets whose database entries are not soft-deleted and for which no migrated user files can be found.

The following menu appears:

```
DAEMON DATABASE BITFILE IDS FOR
WHICH NO FILES CAN BE FOUND
-----
```

There are 64 bitfile IDs in this error class.

7: Examining and Correcting Discrepancies

The following errors were detected in these bitfile IDs.

01. There are 64 MSP/VG daemon database entries that should be soft deleted.

The following actions will be taken to correct these errors.

02. 64 MSP/VG daemon database entry soft-delete flags will be set.

Select:

```
<accept>  Accept the recommended actions. Once accepted, these actions
           will be taken when you select 'apply' on the main menu.
<examine>  Examine or modify the actions to be taken for individual files.
           You may enter the line number of any one of the above subsets of
           files; the default is to examine all the files.
<bfid>     Examine all files and database entries for any bitfile ID you
           specify.
<dump>     Append all information available on each of the bitfile IDs
           listed above to a text file in a format suitable for machine
           processing.
<up>       Return to the previous menu.
```

Please enter your selection:

For each of the error class menus, dmaudit always shows the following:

- Number of BFID sets with errors in this error class
- Summaries of all the individual errors found in the BFID sets in this class
- What actions dmaudit needs to take to fix the individual errors

The messages displayed in these menus depend upon what errors were found on your system. There are five options available on this display. up returns you to the INSPECT menu, and the bfid option functions as described "The bfid Option" on page 41.

The accept Option

The most important option on an error class menu is `accept`. By entering `accept`, you tell `dmaudit` that you are willing to let it take the listed actions to correct the discrepancies in these BFID sets. In this particular case, it means that `dmaudit` will soft-delete 64 daemon database entries. Action is not taken until you enter `apply` at the MAIN menu.

After you have entered `accept`, the format of the menu changes slightly. The `accept` option disappears because it is no longer needed, and in its place you will instead see the following:

```
<cancel>    Cancel your acceptance of the above recommendations.
```

The `cancel` option allows you to change your mind. If you enter `cancel`, it disappears and `accept` reappears, returning the screen to its original state. You can change your mind as many times as you like up until the point at which you enter `apply` on the MAIN menu. At that time, `dmaudit` fixes all BFID sets whose `accept` flag is currently set.

Most of the time, you do not need to examine these BFID sets in more detail and can type `accept` followed by `up` to move on to other error classes. If you do want to go into additional detail, there are additional options available to help you.

The examine Option

The `examine` option allows you to look at individual BFID sets within this error class. The format of its display is very similar to the `bfid` menu. If you enter `examine`, you are prompted for a line number:

```
Line number (<CR> for all)?
```

Each of the errors and each of the actions reported on the error class menu have line numbers on them. If you enter a line number at the `examine` prompt, it examines only information about BFID sets that have the specified error or action associated with them. If you want to examine all BFID sets in this error class, press `ENTER`. Because there are only two BFID sets in this class, assume that you press `ENTER`.

You then see the following:

```
BITFILE ID = 34745a2b000000000000c6c1
```

```
ENTRIES IN USE:
```

7: Examining and Correcting Discrepancies

```
01. no user file -
02. daemon MSP/VG database entry - dev=50331653 ino=3445388 size=692637 uid=11414
    otime=Nov 26 06:04:20 1997 utime=Nov 26 06:04:20 1997 ctime=Nov 26 06:04:20
    1997 dtime=0 name=/NONAME msp/vg=ftp key=<dinesh/34745a2b000000000000c6c1>
        This database entry should be soft deleted.
        The soft delete flag will be set in this entry.
```

Select:

```
<next>    Examine the next bfid in the list
<last>    Examine the last bfid in the list
<mode>    Switch to abbreviated display
<edit>    Edit this bfid
<dump>    Append all information about this bfid to a text file in a format
           suitable for machine processing
<accept>  Accept the changes to this bitfile ID
<up>     Return to the previous menu
```

Enter <CR> to continue:

Please enter your selection:

This display shows an example of a BFID set with an error. Each object in a BFID set that has an error associated with it also has error and action descriptions following it on the screen. In this case, line 2 is followed by the error description and action required to correct the database entry given. It explains that no user file was found with this BFID.

Table 7-3 describes the other options on this menu.

Table 7-3 examine Menu Options

Option	Description
next	Displays the next BFID set within the list of BFID sets you chose to look at. After that BFID is displayed, the <code>prev</code> option appears, allowing you to see previous BFID sets.
last	Goes directly to the last BFID set in the list instead of having to enter <code>next</code> multiple times. The <code>first</code> option, when it appears, allows you to return to the first BFID set in the list.

Option	Description
mode	Toggles back and forth between abbreviated and full versions of the display.
edit	Allows you to make changes to the BFID set to force <code>dmaudit</code> to correct it in some other way. This option is only needed in several very specific cases; it should normally not be used. For more information, see "The edit Option" on page 70.
dump	Dumps all information about the BFID set to a file in machine-readable form. You are prompted for the name of a file to which data should be appended.
accept	Accepts the actions required to correct the BFID set without affecting other BFID sets in the same error class. This option should seldom be needed; normally you will accept all errors in the class at the same time.

The dump Option

The `dump` option is similar to the one on the `bfid` and `examine` menus, but this `dump` outputs all information known about any of the BFID sets in this error class rather than just a single BFID set.

If you enter `dump`, you are prompted for a line number:

```
Line number (<CR> for all)?
```

Enter a line number if you only want to examine BFID sets that have a specified error or action associated with them. If you want to examine all BFID sets in this error class, press `ENTER`. The next prompt asks for a file name:

```
File to which you want the dump appended?
```

Enter the full or relative pathname of a file, or press `ENTER` to cancel.

Returning to the Inspect Menu

Enter `accept` to accept the changes to these two BFID sets, and then enter `up` again to move on to the next error class.

You then see the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
```

```
-----
```

Errors are divided into several major classes. Only those classes that pertain to this snapshot are displayed.

Select:

- <4> Examine files that have correctable errors
- <5> Examine bfid's in the daemon database for which no user files can be found
- <bfid> Examine all files and database entries for any bfid you specify
- <search> Scan file systems for names of all files with errors (very slow)
- <up> Return to the previous menu

Please enter your selection:

Fixing Files with Correctable Errors

The next step in the example is to look at correctable errors. The correctable error class introduces several new options that were not available in the previous error class menu.

If you enter 4, you see the following:

```
FILES WHICH HAVE CORRECTABLE ERRORS
```

```
-----
```

There are 9 bitfile IDs in this error class.

The following errors were detected in these bitfile IDs.

01. There are 9 MSP/VG daemon database entries with an unknown MSP or VG name.
02. There are 9 migrated user files that have no MSP/VG daemon database entries.

The following actions will be taken to correct these errors.

03. 9 MSP/VG daemon database entries will be removed.
04. 9 migrated user files will have incomplete MSP/VG daemon database entries created for their default MSPs/VGs.

05. 9 migrated user files will be remigrated to all MSPs/VGs for which they have incomplete database entries.

Select:

```

<accept>    Accept the recommended actions.  Once accepted, these actions
              will be taken when you select 'apply' on the main menu.
<examine>   Examine or modify the actions to be taken for individual files.
              You may enter the line number of any one of the above subsets
              of files; the default is to examine all the files.
<bfid>      Examine all files and database entries for any bitfile ID you
              specify.
<opt_part>  Recompute all corrective actions assuming that daemon database
              entries with errors can be removed as long as one good copy of
              the file exists elsewhere.
<search>    Scan file systems for the names of all files listed above (very
              slow).
<dump>      Append all information available on each of the bitfile IDs
              listed above to a text file in a format suitable for machine
              processing.
<nlist>     Append all file names for all the files listed above to a text
              file after first scanning the file systems for their names
              (very slow).
<up>       Return to the previous menu.

```

Please enter your selection:

While this display is certainly more complex than the display for the previous error class, its format is essentially the same. The first section on the screen shows those errors that were detected in the BFID sets in this class, and the second section shows what actions `dmaudit` will take to correct those errors.

Recoverable errors tend to be more complex than errors you have seen previously. The following is an example of what a BFID set with recoverable errors looks like:

```
BITFILE ID = 34745a2b0000000000000954b
```

```
ENTRIES IN USE:
```

```
01. dual-state user file - dev 50331652 (daemon), fhandle
    010000000000001888c5f6086f39b65c000e00000000020000000007040f6, uid
```

15948, size 3863854

No MSP/VG daemon database entries exist for this user file.

Incomplete MSP/VG daemon database entries will be created for the default MSPs/VGs associated with this user file.

This file will be remigrated to those MSPs/VGs for which it has incomplete database entries.

The online file is the only available copy of the user file's data. Even though the file is in a dual state, no daemon database entries exist for the file, and therefore no backup copies are available. In order to repair this BFID set, `dmaudit` has to remigrate the user file to all the MSPs and volume groups that have the incomplete entry. Several new options that are available on the recoverable error class menu are described in the following sections.

The `opt_full` and `opt_part` Options

The following BFID set has no MSP or volume group database entry, but the dual-state user file indicates that the online copy of the file is valid. The display is as follows:

```
BITFILE ID = 34745a2b0000000000000954b
```

```
ENTRIES IN USE:
```

```
01. dual-state user file - dev 50331652 (daemon), fhandle
```

```
010000000000001888c5f6086f39b65c000e000000000020000000007040f6, uid
```

```
15948, size 3863854
```

No MSP/VG daemon database entries exist for this user file.

Incomplete MSP/VG daemon database entries will be created for the default MSPs/VGs associated with this user file.

This file will be remigrated to those MSPs/VGs for which it has incomplete database entries.

There are two ways to fix this problem:

- Recreate the MSP and volume group database entries and remigrate the file to all MSPs and volume groups specified in the current configuration.
- Remove the bitfile ID, making the file a regular online file.

As released, `dmaudit` uses the first approach. This can be modified in two ways. If you enter `config` from the `MAIN` menu, one of the displayed options allows you to specify which method `dmaudit` uses as its default. `dmaudit` remembers the new

default from run to run. For more information, see "The invalid Option" on page 86. You can also see the results of using the second approach by entering the `opt_part` option on the current display. When selected, `dmaudit` redisplay the error class menu, replacing the following action:

- 04. 9 migrated user files will have incomplete MSP/VG daemon database entries created for their default MSPs/VGs.
- 05. 9 migrated user files will be remigrated to all MSPs/VGs for which they have incomplete database entries.

The other actions on the display remain the same (except that some lines have different line numbers). If you then enter `examine` to look at the same BFID set previously shown, you see the same change in actions.

After you have entered `opt_part`, it disappears from the menu and is replaced by the `opt_full` option, which toggles back to the first recovery approach. You can toggle back and forth as much as you like. When you enter `accept`, `dmaudit` uses the recovery method currently in force. If you toggle the method after entering `accept`, `dmaudit` performs a cancel on those BFID sets affected by the toggle. You then must accept the changes again.

The search Option

The search option allows you to find the names of all user files associated with BFID sets in this error class. The names of the files are found by recursively scanning the directories of the filesystems in which the files reside. This can take a long time, so select this option only if you truly need to know the file names. If all the names of all the user files in this error class are already known, the search option does not appear.

After the names have been found, they appear in several places, such as on the `examine` menu. The following screen shows the beginning of the display for a BFID used in the previous example:

```
BITFILE ID = 34745a2b0000000000000954b

ENTRIES IN USE:

01. dual-state user file - dev=50331652 (daemon)
   fhandle=01000000000001888c5f6086f39b65c000e000000000020000000007040f6
   size=3863854 uid=15948 nlinks=0
   /dm1/stim/dir1/dmf_tst.04266
   No MSP/VG daemon database entries exist for this user file.
```

The bitfile ID will be removed from this user file.

The BFID (bitfile ID) will be removed from this user file. The pathnames are also output when using the `dump` option, and when using `nlist` (described in the next section). If you need to know the names of user files in more than one error class, it is more efficient to use the `search` option that appears on the `INSPECT` menu instead of using the `search` option repeatedly in each of the error class menus.

The `nlist` Option

The `nlist` option is useful when you want to create a list of the names of user files in a particular error class. `nlist` saves the list in a file whose name you specify. Normally, the `nlist` option looks like the following:

```
<nlist> Append all file names for all the files listed above to a text
file after first scanning the file systems for their names
(very slow).
```

If you have previously used the `search` option to retrieve the names of the files in this error class, the `nlist` option will instead look like the following:

```
<nlist> Append all file names for all the files listed above to a text
file.
```

You are prompted for a line number:

```
Line number (<CR> for all)?
```

You can either enter one of the line numbers appearing on the error class menu to create a list containing a subset of the files in this error class, or you can press `ENTER` to get the names of all user files in this error class. You are then prompted for the name of a file to which the file names should be appended:

```
File to which you want the file names appended?
```


Enter the full or relative pathname of a file, or press ENTER to cancel the operation.
The following display shows a sample file created using the `nlist` option:

```
/dmf1/example/cmp_to_inc  
/dmf1/example/back\e134slash  
/dmf1/example/pipe\e174char  
/dmf1/example/link1  
/dmf1/example/link2  
/dmf1/example/link3  
/dmf1/example/inc_to_cmp  
/dmf1/example/lost_inc_pfile  
/dmf1/example/rmv_inc_pfile
```

The example shows several things that you must be aware of when you use the output from `nlist`, especially when using it as input to some other program:

- If a file has multiple pathnames, `nlist` lists each of the pathnames on consecutive lines. This is done for hard links only; symbolic links to a file are not shown. In the previous example, the pathnames for `link1`, `link2`, and `link3` are all hard links to the same file.
- Because nothing prevents users from creating file names containing line feeds, vertical tabs, and other unprintable characters, `nlist` converts such characters into printable form before adding the pathname to the list. For each backslash, vertical bar, or unprintable character in a file name, `nlist` replaces that character with the string `000`, where the `000` field is a 3-digit octal number that is the ASCII value for that character. In the previous example, file `backslash` contains a backslash and file `pipe|char` contains a vertical bar.

Returning to the Inspect Menu

Assuming that you are done looking at this error class, you would enter `accept` to accept the changes to the BFID sets, and then enter `up` again to return to the INSPECT menu. You then see the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS  
-----
```

```
Errors are divided into several major classes. Only those classes that pertain  
to this snapshot are displayed.
```

```
Select:
```

```
<4>      Examine files that have correctable errors
<5>      Examine bfid's in the daemon database for which no user files
          can be found
<bfid>   Examine all files and database entries for any bfid you specify
<search> Scan file systems for names of all files with errors (very slow)
<up>     Return to the previous menu
```

Please enter your selection:

Because you have now accepted the corrections in both error classes, enter up to return to the MAIN menu where you apply the changes that you have accepted. For more information, see "The apply Option" on page 80.

An xfsdump and xfsrestore can only produce class 4 and class 5 errors. The next few sections show examples of other error classes you might see.

Cleaning Up Files with Unrecoverable Errors

A BFID set contains an *unrecoverable error* if a user inode exists but no valid copy of the user's data can be found. You should never see a BFID set with unrecoverable errors in normal operation. If one does occur, it means that either an administrative error was made, a filesystem error occurred, or that DMF has failed in some way. You should always examine errors in this class to determine why they happened, in order to prevent possible recurrences in the future.

If a partial-state file has an unrecoverable error, part of the file's data may exist on disk.

This section discusses the following:

- "Unrecoverable Errors for Offline Files" on page 56
- "The remove Option for Offline Files" on page 59
- "Unrecoverable Errors for Partial-State Files" on page 60
- "The remove option for Partial-State Files" on page 62

Unrecoverable Errors for Offline Files

To give an example of an unrecoverable error, a test file is migrated offline to alternate media. The `dmddadm` command is then used to manually remove the database entry

for the associated BFID to induce the error. When `dmaudit` is next run, it reports an unrecoverable error. The resulting `INSPECT` menu looks like the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
-----
```

Errors are divided into several major classes. Only those classes that pertain to this snapshot are displayed.

Select:

```
<3>          Examine files whose data cannot be recovered
<bfid>       Examine all files and database entries for any bfid you specify
<search>     Scan file systems for names of all files with errors (very slow)
<up>        Return to the previous menu
```

Please enter your selection:

Entering 3 advances you to the following unrecoverable error class menu:

```
FILES WHOSE DATA CANNOT BE RECOVERED
-----
```

There are 1 bitfile IDs in this error class.

The following errors were detected in these bitfile IDs.

01. There are 1 migrated user files for which no good data copies can be found.

The following actions will be taken to correct these errors.

(No actions will be taken.)

Select:

```
<accept>    Accept the recommended actions. Once accepted, these actions
             will be taken when you select 'apply' on the main menu.
<examine>   Examine or modify the actions to be taken for individual files.
             You may enter the line number of any one of the above subsets of
             files; the default is to examine all the files.
<bfid>      Examine all files and database entries for any bitfile ID you
```

7: Examining and Correcting Discrepancies

specify.

<search> Scan file systems for the names of all files listed above (very slow).

<dump> Append all information available on each of the bitfile IDs listed above to a text file in a format suitable for machine processing.

<nlist> Append all file names for all the files listed above to a text file after first scanning the file systems for their names (very slow).

<remove> After first scanning the file systems for their names, remove files that cannot be recovered (very slow).

<up> Return to the previous menu.

Please enter your selection:

Removing the lost user file is not one of the actions listed in lines 4 and 5. To remove lost files, you must use the remove option as described in "The remove Option for Offline Files" on page 59.

The following is an excerpt from the abbreviated display for the BFID set with an unrecoverable error:

BITFILE ID = 20000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
010000000000001885beb4c83ff373e0000e0000000000420000000000000043, uid 0,
size 4474272
No good data copy can be found for this user file.

Select:

<mode> Switch to full display
<edit> Edit this bfid
<dump> Append all information about this bfid to a text file in a format suitable for machine processing
<accept> Accept the changes to this bitfile ID
<up> Return to the previous menu

Please enter your selection:

The remove Option for Offline Files

The `dmaudit` command does not automatically remove unrecoverable user files when you enter `accept` and `apply`. To actually delete the files, you must select the `remove` option. If you have not used the `search` option, the `remove` option looks like the following:

```
<remove> After first scanning the file systems for their names, remove
files that cannot be recovered (very slow).
```

If you have already used `search` to find the file's pathnames, the `remove` will instead look like the following:

```
<remove> Remove files that cannot be recovered.
```

After you enter the `remove` option, the error class display shows the additional action that the files will be removed:

```
FILES WHOSE DATA CANNOT BE RECOVERED
-----
There are 1 bitfile IDs in this error class.
```

The following errors were detected in these bitfile IDs.

```
01. There are 1 migrated user files for which no good data copies can be
found.
```

The following actions will be taken to correct these errors.

```
02. 1 migrated user files will be removed.
```

You should then enter `accept` to accept the removal of the files. The following is an excerpt from the abbreviated display of the BFID set that contained the unrecoverable error:

```
BITFILE ID = 200000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
   010000000000001885beb4c83ff373e0000e00000000042000000000000043, uid 0,
   size 4474272
```

```
/mig/foo
  No good data copy can be found for this user file.
  This file will be removed.
```

Unrecoverable Errors for Partial-State Files

In the following example of an unrecoverable error, a test file was migrated offline to alternate media. Then a portion of the file was recalled, leaving the file in partial-state state. The `dmdadm` command was then used to manually remove the database entry for the associated BFID to induce the error. When `dmaudit` is next run, it reports an unrecoverable error. The resulting `INSPECT` menu looks like the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
-----
```

Errors are divided into several major classes. Only those classes that pertain to this snapshot are displayed.

Select:

- <10> Examine files in PARTIAL-STATE whose OFFLINE regions cannot be recovered
- <bfid> Examine all files and database entries for any bitfile ID you specify
- <search> Scan file systems for names of all files with errors (very slow)
- <up> Return to the previous menu

Please enter your selection:

Entering 10 advances you to the following unrecoverable error class menu:

```
FILES IN PARTIAL STATE WHOSE OFFLINE REGIONS CANNOT BE RECOVERED
-----
```

There are 1 bitfile IDs in this error class.

The following errors were detected in these bitfile IDs.

- 01. There are 1 migrated user files in PARTIAL-STATE for which no good data copies can be found.

The following actions will be taken to correct these errors.

(No actions will be taken.)

Select:

- <accept> Accept the recommended actions. Once accepted, these actions will be taken when you select 'apply' on the main menu.
- <examine> Examine or modify the actions to be taken for individual files. You may enter the line number of any one of the above subsets of
- <bfid> Examine all files and database entries for any bitfile ID you specify.
- <search> Scan file systems for the names of all files listed above (very slow).
- <dump> Append all information available on each of the bitfile IDs listed above to a text file in a format suitable for machine processing.
- <nlist> Append all file names for all the files listed above to a text file after first scanning the file systems for their names (very slow).
- <up> Return to the previous menu.
- <remove> After first scanning the file systems for their names, remove files that cannot be recovered (very slow). If any of these files are PARTIAL-STATE, the data that is on disk will also be removed.

Please enter your selection:

Removing the lost user file is not one of the actions listed in lines 4 and 5.

Note: This file is a partial-state file and some of its data is resident on disk. Before removing the file, you may want to take steps to save the data that is on disk. To remove lost files, you must use the remove option as described in "The remove option for Partial-State Files" on page 62.

The following is an excerpt from the abbreviated display for the BFID set with this unrecoverable error:

```
BITFILE ID = 3faff5140000000000000d5b
```

ENTRIES IN USE:

01. partial-state user file - dev 748 (daemon), fhandle
01000000000001887a93cd376e04e3600000000409cdd0000002a00000000, uid 0,
size 34816
No good data copy can be found for portions of this user file.

Select:

<mode> Switch to full display
<edit> Edit this bfid
<dump> Append all information about this bfid to a text file in a format
suitable for machine processing
<accept> Accept the changes to this bitfile ID
<up> Return to the previous menu

Please enter your selection:

The remove option for Partial-State Files

The `dmaudit` command does not automatically remove unrecoverable user files when you enter `accept` and `apply`. To actually delete the files, you must select the `remove` option. If you have not used the search option, the `remove` option looks like the following:

<remove> After first scanning the file systems for their names, remove files that cannot be recovered (very slow). If any of these files are `PARTIAL-STATE`, the data that is on disk will also be removed.

If you have already used `search` to find the file's pathnames, the `remove` option will instead look like the following:

<remove> Remove files that cannot be recovered. If any of these files are `PARTIAL-STATE`, the data that is on disk will also be removed.

Note: The file is a partial-state file and some of its data may be resident on disk. Selecting the `remove` option will cause this data to be lost. Before selecting the `remove` option, you may want to take steps to save the data that is on disk.

After you enter the `remove` option, the error class display shows the additional action that the files will be removed:

```
FILES IN PARTIAL STATE WHOSE OFFLINE REGIONS CANNOT BE RECOVERED
-----
```

There are 1 bitfile IDs in this error class.

The following errors were detected in these bitfile IDs.

01. There are 1 migrated user files in PARTIAL-STATE for which no good data copies can be found.

The following actions will be taken to correct these errors.

02. 1 migrated user files will be removed.

You can then enter `accept` to accept the removal of the files. The following is an excerpt from the abbreviated display of the BFID set that contained the unrecoverable error:

```
BITFILE ID = 3faff51400000000000000d5b
```

```
ENTRIES IN USE:
```

01. partial-state user file - dev 748 (daemon), fhandle
010000000000001887a93cd376e04e36000000000409cdd0000002a00000000, uid 0,
size 34816
/dmil/judith/file1
No good data copy can be found for portions of this user file.
This file will be removed.

Cleaning Up Files with the Same BFID but Different Sizes

There are two error classes in which more than one user file can have the same BFID. This section describes those cases in which `dmaudit` is able to use the sizes of the various files to determine which database entries go with which user files.

When the daemon is started, it determines the next available BFID to be allocated by looking in the database to find the highest file number in use for the current database ID. The daemon increments that file number by one and uses it the next time a BFID must be assigned to a user file.

Errors in this class usually happen when an event occurs, just before the daemon stops, that causes the last few database updates to be lost. An example is a system interrupt in which the last updates to the database before the crash never arrive on disk.

The database entries for those BFIDs were not successfully placed in the database. Therefore, when the daemon is restarted, it will not know that those BFIDs are in use and will reallocate them to new user files.

If such an event did occur, you would see the following line as part of the `dmaudit` error report:

There are 1 bfid's in use by more than one file that can be automatically corrected.

Entering `inspect` from the MAIN menu would show the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
```

```
-----  
Errors are divided into several major classes. Only those classes that pertain  
to this snapshot are displayed.
```

```
Select:
```

```
<2>      Examine files with nonunique bitfile IDs that can be  
         automatically corrected  
<5>      Examine bitfile IDs in the daemon database for which no user  
         files can be found  
<bfid>   Examine all files and database entries for any bitfile ID you  
         specify  
<search> Scan file systems for names of all files with errors (very slow)  
<up>     Return to the previous menu
```

```
Please enter your selection:
```

Enter 2 to see the following:

```
FILES WITH NONUNIQUE BITFILE IDS THAT  
CAN BE AUTOMATICALLY CORRECTED
```

```
-----  
There are 1 bitfile IDs in this error class.
```

The following errors were detected in these bitfile IDs.

01. There are 1 migrated user files for which no good data copies can be found.

The following actions will be taken to correct these errors.

02. 1 migrated user files and their database entries will all be given new bitfile IDs.

Select:

- <accept> Accept the recommended actions. Once accepted, these actions will be taken when you select 'apply' on the main menu.
- <examine> Examine or modify the actions to be taken for individual files. You may enter the line number of any one of the above subsets of files; the default is to examine all the files.
- <bfid> Examine all files and database entries for any bitfile ID you specify.
- <search> Scan file systems for the names of all files listed above (very slow).
- <dump> Append all information available on each of the bitfile IDs listed above to a text file in a format suitable for machine processing.
- <nlist> Append all file names for all the files listed above to a text file after first scanning the file systems for their names (very slow).
- <remove> After first scanning the file systems for their names, remove files that cannot be recovered (very slow).
- <up> Return to the previous menu.

Please enter your selection:

All of the menu items on this display have already been explained in the previous sections. What is new in this class is the complexity of some of the errors and the actions needed to resolve them.

The following is an example of an abbreviated display for a BFID set in this error class:

BITFILE ID = 345612ca00000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
010000000000001885beb4c83ff373e0000e0000000001a000000000000048, uid 0,
size 29604
 No good data copy can be found for this user file.

02. off-line user file - dev 33554464 (daemon), fhandle
010000000000001885beb4c83ff373e0000e00000000043000000000000043, uid 0,
size 4474272
 This file and its database entries will all be given a new bitfile ID.

03. daemon MSP/VG <ftpl> database entry - not soft deleted, size 4474272, key
<root/345612ca00000000000000001>

Select:

- <mode> Switch to full display
- <edit> Edit this bfid
- <dump> Append all information about this bfid to a text file in a format
 suitable for machine processing
- <accept> Accept the changes to this bitfile ID
- <up> Return to the previous menu

Please enter your selection:

There is one new, subtle piece of information on this display: the blank line between items 1 and 2, which indicates that more than one file has this BFID, and that the database entry in item 3 goes with the user file in line 2 (because the files have identical sizes), not the user file in item 1. Each group of lines separated from the rest by a blank line will be recovered separately by `dmaudit`.

Because two files have the same BFID, `dmaudit` gives a new BFID to at least one of the files so that there is no overlap of entries in the daemon database. In the previous example, `dmaudit` has decided to leave item 1 as it is. It instead assigns a new BFID to the user file in item 2. It then creates incomplete database entries for the default MSPs and volume groups for UID 0 and migrates the file to those MSPs and volume groups.

The `opt_part` option offers much simpler correction of these errors by not requiring `dmaudit` to leave a file with the same number of MSP or volume group copies as it had before recovery. In this particular case, selecting `opt_part` tells `dmaudit` that it does not need to create missing MSP and volume group entries.

Cleaning Up Multiple Files with the Same BFID and Size

`dmaudit` cannot itself solve errors in which it has found more than one user file has the same BFID and the same size in bytes. There are two main ways this type of error can occur:

- The most common way is when a user file is restored using a `xfsrestore` command. For example, assume that a user owns a migrated file that has multiple pathnames (hard links). Assume further that the user uses the `rm` command to remove one of the pathnames, then asks the administrator to restore the file. Your restore command creates a migrated inode using the pathname supplied by the user even though the original migrated inode is still in the filesystem, pointed to by the remaining hard links. The result is that now two inodes exist in the filesystem with the same BFID and with the same size.
- It is also possible that the assignment of the same BFID to more than one inode is an error such as those described in the previous section, and that it is only a coincidence that the two files have the same size.

In the first case, the two inodes are really meant to contain the same data. In the second case they are not. Furthermore, in the second case, `dmaudit` has no way of knowing which database entry should go with which file.

To resolve these errors, `dmaudit` must rely upon information supplied by you. There are options that you can use to tell `dmaudit` which of the above two cases occurred. You also have the ability to tell `dmaudit` which database entry goes with which file when necessary.

If such an error occurs, the `dmaudit` error report will contain something like the following:

```
DAEMON DATABASE ERROR REPORT
```

```
-----
```

```
There are 1 bfid's in use by more than one file that cannot be corrected without  
additional information from you.
```

If you enter the inspect option from the MAIN menu, you see the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
-----
```

```
Errors are divided into several major classes.  Only those classes that pertain
to this snapshot are displayed.
```

```
Select:
```

- <1> Examine files with nonunique bitfile IDs that cannot be corrected without additional information from you
- <bfid> Examine all files and database entries for any bitfile ID you specify
- <search> Scan file systems for names of all files with errors (very slow)
- <up> Return to the previous menu

```
Please enter your selection:
```

Entering 1 places you at the error class menu:

```
FILES WITH NONUNIQUE BITFILE IDS THAT CANNOT BE
CORRECTED WITHOUT ADDITIONAL INFORMATION FROM YOU
-----
```

```
There are 1 bitfile IDs in this error class.
```

```
The following errors were detected in these bitfile IDs.
```

- 01. There are 2 migrated user files that cannot be cleaned up without additional information from you.

```
The following actions will be taken to correct these errors.
```

```
(No actions will be taken.)
```

```
Select:
```

- <examine> Examine or modify the actions to be taken for individual files. You may enter the line number of any one of the above subsets of files; the default is to examine all the files.
- <bfid> Examine all files and database entries for any bitfile ID you

```

specify.
<search> Scan file systems for the names of all files listed above (very
slow).
<dump> Append all information available on each of the bitfile IDs
listed above to a text file in a format suitable for machine
processing.
<nlist> Append all file names for all the files listed above to a text
file after first scanning the file systems for their names (very
slow).
<up> Return to the previous menu.

```

Please enter your selection:

Because the names of the files are often a good clue as to whether the duplicates were created by your restore command, you may want to use the search option on this menu to collect the names of the files before continuing. When the names have been determined, select the examine option, and press ENTER when you are prompted for a line number. You will want to examine each BFID set in this error class, one at a time.

The following is an example of such a BFID set:

```
BITFILE ID = 200000000000000001
```

```
ENTRIES IN USE:
```

```

01. off-line user file - dev 33554464 (daemon), fhandle
010000000000001885beb4c83ff373e0000e00000000001b0000000000000048, uid 0,
size 4474272
/mig/bar
    You must specify whether this file is a duplicate created by the
    'restore' command.

02. off-line user file - dev 33554464 (daemon), fhandle
010000000000001885beb4c83ff373e0000e0000000000430000000000000043, uid 0,
size 4474272
/mig/foo
    You must specify whether this file is a duplicate created by the
    'restore' command.

03. daemon MSP/VG <ftpl> database entry - soft deleted, size 4474272, key
<root/345612ca00000000000000001>

```

Select:

- <mode> Switch to full display
- <edit> Edit this bfid
- <dump> Append all information about this bfid to a text file in a format suitable for machine processing
- <up> Return to the previous menu

Please enter your selection:

The first step is to decide whether one of the files was created by a restore command. Often the easiest way to determine this is to use the files' UIDs to determine the owner of the files, and then ask the owner if the files were meant to contain the same data.

The edit Option

When you have determined whether the files should contain the same data, you must enter the edit option to pass that information on to dmaudit. The following display appears:

```
EDIT ONE BITFILE ID
-----
BITFILE ID = 20000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
   010000000000001885beb4c83ff373e0000e0000000001b0000000000000048, uid 0,
   size 4474272
   /mig/bar
   You must specify whether this file is a duplicate created by the
   'restore' command.

02. off-line user file - dev 33554464 (daemon), fhandle
   010000000000001885beb4c83ff373e0000e000000000430000000000000043, uid 0,
   size 4474272
   /mig/foo
   You must specify whether this file is a duplicate created by the
   'restore' command.

03. daemon MSP/VG <ftpl> database entry - soft deleted, size 4474272, key
```



```
<root/345612ca0000000000000001>
```

Select:

```
<transfer>    Transfer a database entry to the ownership of a different user
              file
<remove>     Remove objects that you do not want to keep
<nondup>     Indicate that a file is NOT a duplicate created by the
              'restore' command
<duplicate>  Indicate that a file is a duplicate created by the 'restore'
              command
<mode>      Switch to full display
<up>       Return to the previous menu
```

Please enter your selection:

This is the menu from which you give directions to `dmaudit` on how to resolve the error. When you start making changes to the BFID set, you cannot leave this menu until you either accept the changes you have made so far or cancel them. (`accept` and `cancel` will appear when you make your first change.)

When you have given enough information for `dmaudit` to resolve an error, it moves the BFID set to its appropriate new error class. For example, assume that using the `edit` menu you told `dmaudit` that the file on line 1 is a restore duplicate of the file on line 2. Using that new information, `dmaudit` could reclassify the BFID set as having only a recoverable error and would want to move the BFID set to the recoverable error class. When you start to make changes to a BFID set, the `up` option disappears from the menu, keeping you on the `edit` menu until you enter either `accept` or `cancel`. Entering `accept` from this menu means that you accept any reclassification that will take place as a result of your changes. As soon as you enter `accept`, the `up` option appears again, allowing you to leave the display. You can still enter `cancel` at this point if you change your mind. When you select `up` and leave the display, `dmaudit` reclassifies the BFID set and moves it to its appropriate new error class.

If you make some changes to the BFID set, but those changes are insufficient to tell `dmaudit` how to resolve all ambiguities, `dmaudit` will leave the BFID set in the current error class.

When `dmaudit` can solve all the BFID sets in this error class, the error class becomes empty and all the BFID sets are moved to other error classes.

When the BFID sets have been moved, you must go to those error class menus and accept the changes before using apply on the MAIN menu. It is therefore a good idea to resolve BFID sets in this error class first before proceeding to other error classes.

The duplicate Option

Use the duplicate option if you want to indicate to dmaudit that a file is a duplicate created by an xfsrestore command.

You are prompted for the line number of the file that is the duplicate:

```
Line number of duplicate file created by 'restore'?
```

Enter the line number at the prompt. In the example, line 1 is the restored file. You are prompted for the line number of the original file:

```
Line number of file it is a duplicate of?
```

In the example, the original file is on line 2. dmaudit then redisplayes the BFID including errors and actions based upon the new information that you provided.

The example would look like the following:

```
EDIT ONE BITFILE ID
-----
BITFILE ID = 200000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
    010000000000001885beb4c83ff373e0000e00000000043000000000000043, uid 0,
    size 4474272
    /mig/foo

02. off-line user file - dev 33554464 (daemon), fhandle
    010000000000001885beb4c83ff373e0000e0000000001b0000000000000048, uid 0,
    size 4474272
    /mig/bar
    You have specified that this user file is a duplicate created by the
    'restore' command.
    This file will be recalled, and will then be remigrated to the same MSPs/VGs
    using a new bitfile ID.

03. daemon MSP/VG <ftp1> database entry - soft deleted, size 4474272, key
```

```
<root/345612ca0000000000000001>  
  This database entry should not be soft deleted.  
  The soft delete flag will be cleared from this entry.
```

Select:

```
<accept>  Accept the changes you have made  
<remove>  Remove objects that you do not want to keep  
<nondup>   Indicate that a file is NOT a duplicate created by the 'restore'  
           command  
<cancel>  Cancel all changes you have ever made to this bitfile ID  
<mode>    Switch to full display
```

Please enter your selection:

dmaudit now has all the information it needs to resolve the errors in this BFID. If you enter `accept` and then enter `up` to get back to the `INSPECT` menu, you see that the BFID set has moved to the recoverable error class.

The `INSPECT` menu would look like the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS  
-----
```

Errors are divided into several major classes. Only those classes that pertain to this snapshot are displayed.

Select:

```
<4>       Examine files that have correctable errors  
<bfid>    Examine all files and database entries for any bitfile ID you  
           specify  
<up>     Return to the previous menu
```

Please enter your selection:

To finish the cleanup of the BFID set, enter `4`, enter `accept`, then return to the `MAIN` menu (by using the `up` option), and enter `apply`. `dmaudit` then recalls the duplicate file, removes its BFID, and remigrates the file to the same MSPs and volume groups used by the original file. Each file then has its own separate MSP or volume group copies.

The transfer, remove, and nondup Options

The three other options available on the `edit` menu let you make other changes to a BFID set that influences how it is corrected. Although they are not needed very frequently, they are extremely powerful, allowing you to completely control how `dmaudit` resolves its errors.

The `remove` option allows you to specify objects in the BFID set that you want `dmaudit` to discard. The `transfer` option allows you specify to which user file a particular database entry belongs. The `nondup` option allows you to state that a particular user file is not a duplicate created by a `restore` command.

Consider the following example:

```
EDIT ONE BITFILE ID
-----
BITFILE ID = 200000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
   010000000000001885beb4c83ff373e0000e0000000001b0000000000000048, uid 0,
   size 4474272
   /dmf1/example/restore_dir/myfile
   You must specify whether this file is a duplicate created by the
   'restore' command.

02. off-line user file - dev 33554464 (daemon), fhandle
   010000000000001885beb4c83ff373e0000e000000000430000000000000043, uid 0,
   size 4474272
   /dmf1/example/orig_dir/myfile
   You must specify whether this file is a duplicate created by the
   'restore' command.

03. daemon MSP/VG <ftpl> database entry - not soft deleted, size 4474272, key
   <root/345612ca00000000000000001>

Select:
  <transfer>   Transfer a database entry to the ownership of a different user
               file
  <remove>    Remove objects that you do not want to keep
```

```

<nondup>      Indicate that a file is NOT a duplicate created by the
               'restore' command
<duplicate>  Indicate that a file is a duplicate created by the 'restore'
               command
<mode>       Switch to full display
<up>         Return to the previous menu

```

Please enter your selection:

Assume that the file on line 1 is **not** a duplicate of the file on line 2, but is a different file. Assume further that the MSP or volume group database entry is a copy of the file on line 1.

The first step is to indicate that the file in line 1 is not a duplicate. To do this, select the nondup option. You are prompted for the line number of the original user file:

Line number of user file?

After you enter 1, the edit menu looks like the following:

```

EDIT ONE BITFILE ID
-----
BITFILE ID = 200000000000000001

ENTRIES IN USE:
01. off-line user file - dev 33554464 (daemon), fhandle
    010000000000001885beb4c83ff373e0000e0000000001b0000000000000048, uid 0,
    size 4474272
    /dmf1/example/restore_dir/myfile
    No good data copy can be found for this user file.
    You have labeled this migrated user file as being unique.

02. off-line user file - dev 33554464 (daemon), fhandle
    010000000000001885beb4c83ff373e0000e000000000430000000000000043, uid 0,
    size 4474272
    /dmf1/example/orig_dir/myfile
    This file and its database entries will all be given a new bitfile ID.

03. daemon MSP/VG <ftpl> database entry - not soft deleted, size 4474272, key
    <root/345612ca00000000000000001>

```

Select:

7: Examining and Correcting Discrepancies

<accept> Accept the changes you have made
<transfer> Transfer a database entry to the ownership of a different user file
<remove> Remove objects that you do not want to keep
<duplicate> Indicate that a file is a duplicate created by the 'restore' command
<cancel> Cancel all changes you have ever made to this bitfile ID
<mode> Switch to full display

Please enter your selection:

dmaudit now knows that the file on line 1 is not a duplicate. It realizes that two different files have the same BFID, and so that one of its actions is to assign a new BFID to one of the files and its database entries so that there is no longer any overlap.

dmaudit still believes that the MSP or volume group database entry belongs to the file in line 2, so you must now enter transfer to tell dmaudit that the database entry really belongs to the file on line 1. You are asked the line number of the object to reassign:

Line number of object to be transferred?

In this case, the database entry you want to transfer is on line 3, so you enter 3. Next you are asked to which user file the database entry is to be assigned:

User file to transfer it to?

Enter 1. The edit menu then looks like the following:

EDIT ONE BITFILE ID

BITFILE ID = 20000000000000001

ENTRIES IN USE:

01. off-line user file - dev 33554464 (daemon), fhandle
010000000000001885beb4c83ff373e0000e0000000001b000000000000048, uid 0,
size 4474272
/dmf1/example/restore_dir/myfile
You have labeled this migrated user file as being unique.
This file and its database entries will all be given a new bitfile ID.
02. daemon MSP/VG <ftpl> database entry - not soft deleted, size 4474272, key

<root/345612ca0000000000000001>

03. off-line user file - dev 33554464 (daemon), fhandle
 01000000000001885beb4c83ff373e0000e00000000043000000000000043, uid 0,
 size 4474272
 /dmf1/example/orig_dir/myfile
 No good data copy can be found for this user file.

Select:

- <accept> Accept the changes you have made
- <transfer> Transfer a database entry to the ownership of a different user
file
- <remove> Remove objects that you do not want to keep
- <duplicate> Indicate that a file is a duplicate created by the 'restore'
command
- <cancel> Cancel all changes you have ever made to this bitfile ID
- <mode> Switch to full display

Please enter your selection:

Note: A side effect of moving objects around is that many objects may end up with different line numbers after the transfer. Be aware of this when making subsequent requests.

The final step is to remove the unwanted user file that has no data. However, if in another case you wanted to remove an entry, enter remove, which prompts you for the line number of the object to remove:

Line number of object to be removed?

Enter the line number for the user file with no data. The edit menu would then look like the following:

```
EDIT ONE BITFILE ID
-----
BITFILE ID = 20000000000000001
```

ENTRIES IN USE:

7: Examining and Correcting Discrepancies

01. off-line user file - dev 33554464 (daemon), fhandle
01000000000001885beb4c83ff373e0000e00000000001b0000000000000048, uid 0,
size 4474272
/mig/bar
You have labeled this migrated user file as being unique.
02. daemon MSP/VG <ftpl> database entry - not soft deleted, size 4474272, key
<root/345612ca00000000000000001>

DISCARDED ENTRIES:

03. off-line user file - dev 33554464 (daemon), fhandle
01000000000001885beb4c83ff373e0000e0000000000430000000000000043, uid 0,
size 4474272
/mig/foo
You have chosen to remove this migrated user file.
This file will be removed.

Select:

- <accept> Accept the changes you have made
- <transfer> Transfer a database entry to the ownership of a different user
file
- <remove> Remove objects that you do not want to keep
- <nondup> Indicate that a file is NOT a duplicate created by the
'restore' command
- <duplicate> Indicate that a file is a duplicate created by the 'restore'
command
- <cancel> Cancel all changes you have ever made to this bitfile ID
- <mode> Switch to full display

Please enter your selection:

BFID Sets That Cannot Be Immediately Corrected

Occasionally, dmaudit encounters a BFID set containing errors that is being migrated or unmigrated by someone at the time of the snapshot. In this case, dmaudit is unable to correct the errors because it knows its image of the BFID set will be inaccurate once the migrate or unmigrate request completes.

Sometimes a BFID set with errors is not active at the time of the snapshot, but when `dmaudit` searches for the pathnames of files in the BFID set, it finds that their inodes have changed since the snapshot was taken. In this case, `dmaudit` knows that its information is out-of-date.

In such cases, `dmaudit` still reports the errors that were detected, but the BFID set is placed in a special class in which it can be examined but cannot be corrected. The daemon error report shows something like the following:

There are 1 bfid's whose errors cannot be corrected until a new snapshot is taken.

The INSPECT menu shows the following:

```
INSPECT FILE SYSTEM AND DATABASE ERRORS
-----
```

```
Errors are divided into several major classes. Only those classes that pertain
to this snapshot are displayed.
```

```
Select:
```

```
<6>          Examine bfid's that have changed since the snapshot was taken
<bfid>       Examine all files and database entries for any bfid you specify
<search>     Scan file systems for names of all files with errors (very slow)
<up>        Return to the previous menu
```

```
Please enter your selection:
```

If you enter the error class menu number, you see the following:

```
BITFILE IDS THAT HAVE CHANGED
SINCE THE SNAPSHOT WAS TAKEN
-----
```

```
There are 1 bitfile IDs in this error class.
```

The following errors were detected in these bitfile IDs.

01. There are 1 migrated user files for which no good data copies can be found.
02. There are 1 user files in an unexpected DMF state.

The following actions will be taken to correct these errors.

03. 1 migrated user files and their database entries will all be given new bitfile IDs.

Select:

- <examine> Examine or modify the actions to be taken for individual files. You may enter the line number of any one of the above subsets of files; the default is to examine all the files.
- <bfid> Examine all files and database entries for any bitfile ID you specify.
- <search> Scan file systems for the names of all files listed above (very slow).
- <dump> Append all information available on each of the bitfile IDs listed above to a text file in a format suitable for machine processing.
- <nlist> Append all file names for all the files listed above to a text file after first scanning the file systems for their names (very slow).
- <up> Return to the previous menu.

Please enter your selection:

All the normal options are available with the notable exception of `accept`. `accept` does not appear because `dmaudit` knows its information is out-of-date.

To correct such errors, you must wait until the daemon activity for these BFID sets has completed and then take a new snapshot. Usually, you can use the `dmdidle` command to force outstanding migrate requests to be completed.

The apply Option

When you have used the `accept` option on each of the error class menus to select the errors you want to correct, enter `up` repeatedly until you arrive at the MAIN menu and the `apply` option:

MAIN MENU

Select:

- <apply> Apply all the changes you have accepted
- <inspect> Inspect and correct file system and database errors
- <report> Reprint status report for the current snapshot

```
<verifymsp> Check the dmatls tape library server databases against the
              daemon databases
<snapshot>  Take a snapshot and report status of file systems and databases
<free>      Release all file space used by the current snapshot
<config>    Examine or modify configuration information
<quit>      Quit
```

Please enter your selection:

When you enter `apply`, `dmaudit` corrects all errors that you have selected with the `accept` options. You can do this one of two ways:

- If the actions listed in the error class displays indicate that relatively little daemon, tape, and disk activity is required to correct the errors, you can enter `apply` at the prompt and wait. If `dmaudit` has difficulty correcting any of the errors, it issues messages to your screen. When the `MAIN` menu next appears, the errors that can be corrected will have been resolved.
- If you anticipate that significant activity will occur while recovering the errors, you may want to wait until off-peak hours. At that time you can select `apply` interactively, or you can run `dmaudit` in batch mode. For example, you could log off while `dmaudit` continued to correct errors if you enter the following (using the Bourne shell):

```
nohup dmaudit apply 2>errs &
```

When `dmaudit` finishes executing, any problems that occurred are listed in the `errs` file.

`dmaudit` usually has no difficulty resolving errors. The only time problems occur is if the BFID set has changed since the snapshot was taken. This could happen if the owner of a file has recently tried to recall or remove it. `dmaudit` always first verifies that a BFID set has not changed since the snapshot was taken before beginning to correct its errors. If the BFID set has changed, `dmaudit` reports that fact and continues with the next BFID set. All errors that cannot be corrected are left as is; you must take a fresh snapshot to resolve the BFID sets that changed.

The `dmaudit` error report is updated dynamically as errors are corrected. For example, if you enter `report` from the `MAIN` menu after all errors have been corrected, you see the following:

```
No errors were discovered comparing the file systems against the daemon database.
```

If you see that files will be recalled, ensure that there is sufficient room in the filesystems.

Changing the Configuration

This chapter describes the following `dmaudit` configuration options that are available for you to examine or change:

- "The `workdir` Option" on page 84
- "The `fileSYS` Option" on page 85
- "The `invalid` Option" on page 86
- "The `dmfconf` Option" on page 88
- "The `remove` Option" on page 89

`dmaudit` stores the settings of these options in its configuration file so that they can be used in all future runs. All the options are reached from the `config` option on the MAIN menu.

This chapter assumes that you are executing `dmaudit` interactively, and that you are currently positioned at the MAIN menu. If you are not running `dmaudit`, enter the following:

```
dmaudit
```

If you are running `dmaudit` but are not at the MAIN menu, enter `up` as many times as necessary until you arrive there.

The MAIN menu should look something like the following:

```
MAIN MENU
-----
```

```
Select:
```

```
<snapshot>  Take a snapshot and report status of file systems and databases
<config>    Examine or modify configuration information
<quit>      Quit
```

```
Please enter your selection:
```

Your menu may display more items than in the preceding example if you have already taken a snapshot. You can still enter the CONFIGURATION menu and examine

or change configuration options, but some of those options may cause your snapshot to be discarded. All such options ask for confirmation before discarding information.

Enter `config` to select the CONFIGURATION menu. `dmaudit` responds with the following:

```
CONFIGURATION MENU
-----
```

You may examine or change any of the configuration options below. Changes take effect when you return to the Main menu.

Select:

```
<workdir>  Change the location of the working directory
<filesystem> View or edit the list of file systems that will be scanned
<invalid>  Choose whether invalid MSP and/or VG copies are to be removed
<dmfconf>  Display the current DMF configuration
<remove>   Remove all configuration and snapshot information
<up>      Return to the main menu
```

CONFIGURATION menu items are described in more detail in the following sections, which explain why you would want to use an option and show the necessary steps.

When you are done examining or changing the configuration, enter `up` to return to the MAIN menu. All configuration changes you have made take effect at that time.

The `workdir` Option

The `workdir` option on the CONFIGURATION menu allows you to see where the working directory used by `dmaudit` is currently located and gives you the option of moving the directory to a new location. Because the contents of the previous working directory are not preserved during the move, you should only change the location of the directory if you do not intend to use any current snapshot.

To see where the working directory is, type `workdir` at the CONFIGURATION menu prompt.

If you do not have a current snapshot, `dmaudit` responds with the following:

```
Currently the working directory 'working_dir' exists within directory
'/usr/tmp'.
```

Do you wish to change the location of the working directory (y/n)?

If you do have a current snapshot, `dmaudit` instead responds with the following:

```
Currently the working directory 'working_dir' exists within directory
'/usr/tmp'.
```

Changing the location of the working directory at this point will cause your current snapshot to be discarded. Do you still wish to continue (y/n)?

If you answer `n` to either question, `dmaudit` returns immediately to the CONFIGURATION menu. If you reply `y`, `dmaudit` responds with the following:

```
Please enter the full path name of an existing directory in which subdirectory
'working_dir' can be created (<CR> to quit):
```

Enter the full pathname of the new directory you want to use. `dmaudit` verifies that the directory you specified exists. If it does not, or if any other errors are detected, `dmaudit` describes the nature of the error and then prompts you for a new directory pathname.

If at any time you decide not to change the directory, press `ENTER` to return to the CONFIGURATION menu. The original working directory and its contents remain unchanged.

After you have entered a valid pathname, `dmaudit` creates the new working directory, removes the previous directory, and returns you to the CONFIGURATION menu.

The `filesys` Option

The `filesys` option on the CONFIGURATION menu is useful when you want to see which filesystems `dmaudit` searches for migrated files, or when you want to add to or subtract from that list. If you enter `filesys`, you see the following:

```
SELECT FILE SYSTEMS TO BE SEARCHED
-----
```

The following menu allows you to look at and modify the list of file systems to be scanned by this program when it searches for every migrated file in the system.

Select:

```
<view>  View the current file system scan list
<edit>  Edit a fresh copy of the file system scan list
<up>    Return to the previous menu
```

This menu is very similar to the filesystem selection menu displayed during the initial configuration of `dmaudit`, as shown in "Specifying Filesystems to Scan" on page 24. The `view` option lets you see the current filesystem list, and `edit` allows you to change the list.

Unlike the initial configuration display, this menu does not have an `accept` option. Instead, any changes to the filesystem list take effect when you next return to the MAIN menu. `dmaudit` must discard any current snapshot when the filesystem list is modified; if you select `edit` while a snapshot exists, you are asked to confirm your actions:

```
Editing the file system scan list at this point will cause your current
snapshot to be discarded. Do you still wish to continue (y/n)?
```

A reply of `n` leaves the current filesystem scan list unchanged.

Note: If you do decide to use `edit`, be very careful to ensure that all filesystems that contain migrated files appear in your new list.

The invalid Option

"The `opt_full` and `opt_part` Options" on page 52 gives an example of a particular kind of error that can be corrected in one of two ways. The example shows a BFID set with no database entry but the dual-state user file indicates that the online copy of the file is valid.

The BFID set that has one or more invalid database entries (or has no database entry at all) and has a valid copy (online or offline) of the file can be corrected by doing one of the following:

- Recalling the user file (if no valid online copy exists) and then remigrating it to all MSPs and volume groups (VGs) specified in the configuration. Use this method if it is mandatory that a particular number of copies of a file always exists.
- Removing any incomplete MSP and volume group database entries. If no valid MSP or volume group database entries then exist, the BFID on the user file is removed, making it a regular file.

dmaudit uses the first method by default, but you have the option of changing that default value using the invalid selection on the CONFIGURATION menu.

If you select invalid, the following menu appears:

```
SELECT WHETHER INVALID/MISSING MSP AND/OR VG COPIES CAN BE REPLACED
```

```
-----
```

If this program discovers a file that has

- (a) one or more invalid database entries or
- (b) no database entries at all

and also has

- (a) at least one good database entry and/or
- (b) a good online copy,

it uses one of the following methods to clean up the discrepancies.

1. Replace the invalid/missing database entries by recalling the file (or using a good online copy) and remigrating it to all MSPs and/or VGs whose original copies were invalid/missing.

This method is used by those sites that have a requirement that any migrated file must have back-up copies on a particular set of MSPs and/or VGs.

2. Remove the invalid database entries, leaving any good MSP and/or VG copies as the only back-up copies of the file. If no valid database entries remain, the file becomes a regular file.

This method is acceptable to some sites, and has the advantage of minimizing the amount of MSP/VG activity (such as tape mounts) and file system activity (such as file recalls) generated during the recovery phase of this program.

Currently, invalid MSP/VG daemon database entries will be replaced during database recovery.

Select:

- <1> Replace invalid (or missing) MSP and/or VG daemon database entries
- <2> Remove invalid MSP and/or VG daemon database entries

<up> Return to the previous menu

Please enter your selection:

The menu shows you the default method that will be used in all future runs of `dmaudit` and gives you the option of changing that default. Enter `up` if you do not want to change the default; otherwise, select the number corresponding to the method you want to use.

Even if you do change the default value, you still will be shown the `opt_part` and `opt_full` options when such errors occur, so you still will always have the option of overriding the default value at that time.

The `dmfconf` Option

The `dmfconf` option on the `CONFIGURATION` menu displays information from the DMF configuration file that defines the configuration of DMF at your site.

`dmaudit` collects and saves this information every time a new snapshot is taken, so the `dmfconf` display is actually showing you the DMF configuration at the time the last snapshot was taken.

The following example illustrates the format of the output of the `dmfconf` selection. After the information has been displayed, `dmaudit` automatically returns to the `CONFIGURATION` menu:

```
DATA MIGRATION CONFIGURATION
-----

Data migration home directory:      /dmfhome
Data migration binaries directory:  /usr/sbin

Server name:                        daemon
Server home directory:              /dmfhome/daemon
Server spool directory:              /dmfspool/daemon

Data migration daemon process ID:   1366252

MSP name    MSP type
-----
```

```

ftpl      dmftpmsp

Name      Type
-----
ls        dmatls
vgla     volume group   in Library Server ls

```

CONFIGURATION MENU

Enter <CR> to continue: **ENTER**

You may examine or change any of the configuration options below. Changes take effect when you return to the main menu.

Select:

```

<workdir> Change the location of the working directory
<fileys>  View or edit the list of file systems that will be scanned
<invalid> Choose whether invalid MSP and/or VG copies are to be removed
<dmfconf> Display the current DMF configuration
<remove>  Remove all configuration and snapshot information
<up>     Return to the main menu

```

Please enter your selection:

The remove Option

If for some reason you want to remove your `dmaudit` working directory and configuration file and begin again as if you had never used `dmaudit`, you can use the `remove` option on the CONFIGURATION menu. The `remove` option gives you the ability to remove both the working directory and the configuration file used by `dmaudit`. If you enter this option, you are asked to confirm your actions:

```
All configuration and snapshot information maintained by this program will be
removed. Do you still wish to continue (y/n)?
```

If you answer `y`, all `dmaudit` configuration and snapshot information is removed, and `dmaudit` then silently exits. The next time that you run `dmaudit` you must reconfigure it.

dump Option Output

This appendix describes the fields that appear in the `dmaudit` command `dump` option output. `dmaudit` dumps information about a BFID set as a series of consecutive lines. The BFID is the second field. The first field of each line details the type of data that line contains.

Each type of output produced by `dmaudit` is shown. An example of each is given, and each field within a line is described. `dmaudit` prints each of the data records defined below as one long line. However, long lines in the examples have been broken into multiple lines for readability.

ufile_data

The following is an example of `ufile_data` output:

```
ufile_data|345612ca0000000000000001|  
010000000000001885beb4c83ff373e000e0000000001b000000000000048|0|4474283|1|2|0|2
```

Table A-1 `ufile_data` Output

Field	Description
1	<code>ufile_data</code>
2	Bitfile ID (BFID)
3	File handle (fhandle)
4	File user ID
5	File size (in bytes)
6	File link count

Field	Description
7	File state. Valid values are: <ul style="list-style-type: none"> 0 (a regular file that has data but no BFID) 2 (a dual-state file that has data and a BFID) 3 (an offline file that has no data but has a BFID) 5 (file is exempt from migration) 6 (a partial-state file that has a BFID) 7 (missing file; just a place-holder)
8	Name list offset. A 0 (zero) indicates that the names of the file are not known; a nonzero value indicates that the names are known.
9	Region state. The state of all regions in the file. The following values may be joined together with a logical OR: <ul style="list-style-type: none"> 1 (migrating) 2 (dual-state) 4 (offline) 8 (unmigrating)

ufile_name

If present, these lines immediately follow the ufile_data line to which they apply. There is one line for each path name to a file. Backslash, vertical bar, and nonprintable characters are expanded using the format \000, where 000 is the 3-digit octal ASCII value for that character.

The following is an example of ufile_name output:

```
ufile_name|345612ca0000000000000001|/mig/bar
```

Table A-2 ufile_name Output

Field	Description
1	ufile_name
2	BFID
3	Path name: unprintable characters expanded

ufile_error

If present, these lines follow the ufile_data line to which they apply. The following is an example of ufile_error output:

```
ufile_error|345612ca0000000000000001|No MSP/VG daemon database entries  
exist for this user file.
```

Table A-3 ufile_error Output

Field	Description
1	ufile_error
2	BFID
3	Error message text

ufile_action

If present, these lines follow the ufile_data line to which they apply. The following is an example of ufile_action output:

```
ufile_action|345612ca0000000000000001|The bitfile ID will be removed  
from this user file.
```

Table A-4 ufile_action Output

Field	Description
1	ufile_action
2	BFID
3	Action message text

mdmdb_data

The following is an example of mddb_data output:

```
mdmdb_data|345612ca0000000000000001|33554464|67|4474272|Dec 4 10:11:01 1997| Dec  
4 10:11:01 1997|Dec 4 10:11:01 1997|0|0|foo|ftp1|root/345612ca0000000000000001
```

Table A-5 mddb_data Output

Field	Description
1	mbmbd_data
2	BFID
3	origdv
4	origino
5	origsz
6	otime
7	utime
8	ctime
9	delflag
10	userid
11	ofilenm
12	proc
13	path

mdmdb_error

If present, these lines follow the mddb_data line to which they apply. The following is an example of mddb_error output:

```
mdmdb_error|345612ca0000000000000001|This database entry should not  
be soft deleted.
```


Table A-6 mdmdb_error Output

Field	Description
1	mdmdb_error
2	BFID
3	Error message text

mdmdb_action

If present, these lines follow the mdmdb_data line to which they apply. The following is an example of mdmdb_action output:

```
mdmdb_action|345612ca0000000000000001|The soft delete flag will be cleared
from this entry.
```

Table A-7 mdmdb_action Output

Field	Description
0	mdmdb_action
1	BFID
2	Action message text

Glossary

active database entry

A valid daemon database entry. See also *soft-deleted database entry* and *hard-deletion of database entries*.

alternate media

The media onto which migrated data blocks are stored, usually tapes.

automated space management

The combination of utilities that allows DMF to maintain a specified level of free space on a filesystem through automatic file migration.

base object

The configuration object that defines path name and file size parameters necessary for DMF operation.

BFID

See *bit-file identifier (BFID)*

bit-file identifier (BFID)

A unique identifier, assigned to each file during the migration process, that links a migrated file to its data on alternate media.

BFID set

The collection of database entries and the user file associated with a particular BFID.

BFID-set state

The sum of the states of the components that comprise a BFID set: the file state of any user file and the state of any database entries (incomplete, complete, soft-deleted, or active).

block

Physical unit of I/O to and from media, usually tape. The size of a block is determined by the type of device being written. A tape block is accompanied by a header identifying the chunk number, zone number, and its position within the chunk.

candidate list

A list that contains an entry for each file in a file system eligible for migration, or for a file or range of a file eligible for making offline, ordered from largest file weight (first to be migrated) to smallest. This list is generated and used internally by `dmfsmon`. The `dmSCANFS` command prints similar file status information to standard output.

CAT records

The catalog (CAT) records in the LS database that track which migrated files reside on which tape volumes.

chunk

That portion of a user file that fits on the current media (tape) volume. Most small files are written as single chunks. When a migrated file cannot fit onto a single volume, the file is split into chunks.

complete MSP or volume group daemon-database entry

An entry in the daemon database whose `path` field contains a key returned by its MSP or volume group, indicating that the MSP or volume group maintains a valid copy of the user file.

compression

The mechanism provided by the volume group for copying active data from volumes that contain largely obsolete data to volumes that contain mostly active data. This process is also known as *volume merging* or *tape merging*.

configuration object

A series of parameter definitions in the DMF configuration file that controls the way DMF operates. By changing the parameters associated with objects, you can modify the behavior of DMF.

configuration parameter

A string in the DMF configuration file that defines a part of a configuration object. By changing the values associated with these parameters, you can modify the behavior of DMF. The parameter serves as the name of the line. Some parameters are reserved words, some are supplied by the site.

daemon database

A database maintained by the DMF daemon. This database contains such information as the BFID, the MSP or volume group name, and MSP or volume group key for each copy of a migrated file.

daemon object

The configuration object that defines parameters necessary for `dmfdaemon` operation.

data-pointer area

The portion of the inode that points to the file's data blocks.

device object

The configuration objects that define parameters for the DMF backup scripts' use of tape devices other than those defined by a drive group.

direct-access storage device (DASD)

An IBM disk drive.

disk cache manager (DCM)

The feature that lets you configure the disk MSP to manage data on secondary storage, allowing you to further migrate the data to tape as needed.

DG

See *drive group*.

DMF state

See *file state*.

drive group (DG)

One of the components of an LS. The drive group is responsible for the management of a group of interchangeable tape drives located in the tape library. These drives can be used by multiple volume groups and by non-DMF processes, such as backups and interactive users. However, in the latter cases, the drive group has no management involvement; the mounting service (TMF or OpenVault) is responsible for ensuring that these possibly competing uses of the tape drives do not interfere with each other. The main task of the drive group is to monitor tape I/O for errors, attempt to classify them as volume, drive, or mounting service problems, and to take preventive action.

dual-state file

A file whose data resides both online and offline.

dual-state filesystems

Those filesystems that have the necessary inode space to support dual-state files.

fhandle

See *file handle*.

file

An inode and its associated data blocks; an empty file has an inode but no data blocks.

file handle

The DMAPI identification for a file. You can use the `dmscanfs`, `dmattr`, and `dmfind` commands to find file handles.

file state

The migration state of a file as indicated by the `dmattr` command. A file can be regular (not migrated), migrating, dual-state, offline, partial-state, unmigrating, never-migrated, or have an invalid DMF state.

freed file

A user file that has been migrated and whose data blocks have been released.

freed BFID-set state

A BFID-set state that consists of an offline user file and one or more active, complete MSP and/or volume group database entries.

fully backed up file

A file that has one or more complete offline copies and no pending or incomplete offline copies.

fully migrated BFID-set state

A BFID-set state that consists of one or more active complete MSP and/or volume group database entries, no incomplete database entries, and either a migrated user file or an offline user file.

hard-deletion of database entries

The administrative action that removes an MSP or volume group database entry from the daemon database and discards the MSP or volume group copy. See also *active database entry* and *soft-deleted database entry*.

inode

The portion of a file that contains the BFID, the state field, and the data pointers.

incomplete MSP or volume group daemon-database entry

An entry in the daemon database for an MSP or volume group that has not finished copying the data, and therefore has not yet returned a key. The `path` field in the database entry is NULL.

incompletely migrated BFID-set state

A BFID-set state that consists of a migrating user file; one or more active, incomplete MSP or volume group database entries; and possibly one or more complete active MSP or volume group database entries.

incompletely migrated file

A file that has begun the migration process, but for which one or more copies on alternate media have not yet been made.

library server (LS)

The daemon-like process by which data blocks are copied onto tape and which maintains the location of the migrated data. Each LS has an associated catalog (CAT) and volume (VOL) database. An LS can be configured to contain one or more drive groups. Each drive group contains one or more volume groups. A volume group is responsible for copying data blocks onto alternate media. A volume group is capable of managing a single copy of a user file.

LS

See *library server*

media-specific process (MSP)

The daemon-like process by which data blocks are copied onto alternate media, and which assigns keys to identify the location of the migrated data.

merging

The mechanism provided by the LS for copying active data from volumes that contain largely obsolete data to volumes that contain mostly active data. This process is also known as *volume merging* or *tape merging*.

migrated file

A file that has a BFID and whose offline copies (or copy) are completed. Migrated files can be *dual-state* or *offline*.

migrating file

A file that has a BFID but whose offline copies (or copy) are in progress.

MSP

See *media-specific process (MSP)*.

MSP or volume group database entry

The daemon database entry for a file that contains the path or key that is used to inform a particular MSP or volume group where to locate the copy of the file's data.

MSP objects

The configuration objects that define parameters necessary for that MSP's operation. There is one MSP object for each MSP.

nonmigrated file

A file that does not have a BFID or any offline copies. See *regular file*.

offline file

A file whose inode contains a BFID but whose disk blocks have been removed. The file's data exists elsewhere in copies on alternate media.

offline pointer

In MSP and volume group processing, a character string that the MSP or volume group returns to the daemon to indicate how a file is to be retrieved.

orphan chunks

Unused chunks in the library server (LS) catalog (CAT) database entries resulting from the removal of migrated files.

orphan database entries

Unused database entries resulting from the removal of migrated files during a period in which the DMF daemon is not running.

parameter

See *configuration parameter*.

partial-file online retention

A partial-state file feature capability that allows you to keep a specific region of a file online while freeing the rest of it (for example, if you wanted to keep just the beginning of a file online). See also *partial-state file*.

partial-state

A file that has more than one region. DMF allows for up to four distinct file regions. The file's inode contains a BFID and a complete MSP or volume group `dmfdaemon`

database entry exists for the file. It contains a mixture of DUALSTATE, OFFLINE, or UNMIGRATING regions. See also *region*.

partial-file recall

A partial-state file feature capability that allows you to recall a specific region of a file without recalling the entire file. For more information, see the `dmput(1)` and `dmget(1)` man pages. See also *partial-state file*.

policy objects

The configuration objects that specify parameters to determine MSP or volume group selection, automated space management policies, and/or file weight calculations in automatic space management.

recall

To request that a migrated file's data be moved back (unmigrated) onto the filesystem disk, either by explicitly entering the `dmget(1)` command or by executing another command that will open the file, such as the `vi(1)` command.

region

A contiguous range of bytes that have the same residency state.

region

A portion of a partial-state file, which can itself have a state of OFFLINE, DUALSTATE, or UNMIGRATING.

regular file

DMF considers a regular file to be one with no BFID and no offline copies.

site-defined policies

A site-specific library of C++ functions that DMF will consult when making decisions about its operation.

snapshot

The information about all BFID sets that is collected and analyzed by `dmaudit`. The snapshot analysis is available from the `report` function.

soft-deleted database entry

A daemon database entry for which the MSP or volume group copy of the data is no longer valid. Data remains on the alternate media until the database entry is hard-deleted. See also *active database entry* and *hard-deletion of database entries*.

sparse tape

A tape containing only a small amount of active information.

special file

UNIX special files are never migrated by DMF.

state field

The field in the inode that shows the current migration state of a file.

tape block

See *block*.

tape chunk

See *chunk*.

unmigratable file

A file that the daemon will never select as a migration candidate.

unmigrate

See *recall*.

unmigration file

A special file created in the user file's filesystem by the DMF daemon when a file is recalled. The unmigration file holds the data pointers until the MSP or volume group process successfully copies the file's data back onto the filesystem disk from the alternate media. When the copy is complete, the daemon moves the data pointers from the unmigration file inode back into the user file's inode.

unmigration directory

A directory in which unmigration files are stored.

VG

See *volume group*

voided BFID-set state

A BFID-set state that consists of one or more soft-deleted daemon database entries, either incomplete or complete. There is no user file.

voiding the BFID

The process of removing the BFID from the user file inode and soft-deleting all associated database entries.

VOL records

The volume (VOL) records in the library server (LS) database that contain information about each tape volume that exists in the pool of tapes used by the LS.

volume group (VG)

One of the components of an LS. A volume group is responsible for copying data blocks onto alternate media. Each volume group contains a pool of tapes, all of the same media type, capable of managing single copies of user files. Multiple copies of the same user files require the use of multiple volume groups. See also *library server (LS)*.

volume merging

The mechanism provided by the LS for copying active data from volumes that contain largely obsolete data to volumes that contain mostly active data.

zone

A logical grouping of chunks. Zones are separated by file marks and are the smallest block-addressable unit on the tape volume. The target size of a zone is configurable by media type.

Index

A

accept menu option, 25, 47
apply menu option, 80

B

batch mode
 taking snapshots, 35
BFID
 defined, 5
 set, 6
 state change summary, 17
 use, 10
bfid, 1
BFID database field, 13
bfid menu option, 41
BFID set
 fixing when no user files exist, 45
 states, 6
 diagram, 6
 with migrating files, 79
bit file identifier
 See "BFID", 1

C

COLUMNS environment variable
 dmaudit interactive use, 21
complete database entry
 daemon database, 12
config menu option, 28
configuration
 working directory, 23
correctable errors

 fixing files with, 50
 correcting errors, 37
 cron, 3
 ctime database field, 14

D

daemon database
 contents, 12
 database entry
 complete, 12
 incomplete, 12
 soft-deleted, 12
 fields, 13
daemon internals, 5
data file
 daemon database, 12
database fields
 descriptions, 44
dbrec.keys file, 12
dbrec.path file, 12
delflag database field, 14
discrepancies
 correcting, 37
 detecting with snapshots, 29
disk cache manager (DCM), 3
display size
 dmaudit interactive use, 21
dmaudit, 1
dmaudit_dir/checkpoint file, 23
dmfconf menu option, 88
dual-state file
 defined, 9
DUALSTATE field state, 11
dump menu option, 49
dump option

- output, 91
- duplicate menu option, 72

E

- edit menu option, 70
- editors
 - dmaudit interactive use, 21
- environment variables
 - dmaudit interactive use, 21
- error classes, 37
- error detection, 19
- examine menu option, 47
- example report
 - detecting discrepancies, 34

F

- file state field, 11
- fileys menu option, 85
- filesystems
 - menu for selection, 24
 - removing from dmaudit, 25
 - specifying for scanning, 24
- free menu option, 33
- freeing disk space, 33

H

- hard-deleted, 2

I

- incomplete database entry
 - daemon database, 12
- incompletely migrated state file
 - defined, 9
- index file

- daemon database, 12
- initial configuration, 23
- inode resident data, 10
- INSPECT menu, 39
- inspect menu option, 32
- interactive dmaudit, 21
- invalid menu option, 86

L

- LINES environment variable
 - dmaudit interactive use, 21

M

- Main menu, 30
- mdmdb_action output, 95
- mdmdb_data output, 94
- mdmdb_error output, 94
- menu options
 - accept, 25, 47
 - apply, 80
 - bfid, 41
 - config, 28
 - dmfconf, 88
 - dump, 49
 - duplicate, 72
 - edit, 70
 - examine, 47
 - fileys, 85
 - free, 33
 - inspect, 32
 - invalid, 86
 - nlist, 54
 - nondup, 74
 - opt_full, 52
 - opt_part, 52
 - remove, 74, 89
 - offline files, 59

- partial-state files, 62
- report, 32
- search, 44, 53
- snapshot, 30
- transfer, 74
- up, 30
- verifymsp, 33
- view, 25
- workdir, 84
- menus
 - configuration, 84
 - filesystem selection, 24
 - inspect, 39
 - main, 30
- MIGRATING field state, 11
- migrating file
 - defined, 9
- migrating files
 - correcting BFID sets with, 79
- migration
 - errors during, 15
- migration life-cycle, 6

N

- nlist menu option, 54
- NOMIGR field state, 11
- nondup menu option, 74

O

- OFFLINE field state, 11
- offline file
 - defined, 9
- offline files
 - unrecoverable errors and, 56
- offline-state file
 - defined, 9
- ofilenm database field, 14
- opaque, 5

- opt_full menu option, 52
- opt_part menu option, 52
- origdv database field, 13
- origino database field, 14
- origsz database field, 14
- otime database field, 14

P

- page-wait mechanism, 21
- partial-state file, 10, 17, 38
 - defined, 9
 - region state changes, 6
 - remove option, 62
 - unrecoverable errors, 60
- PARTIALSTATE field state, 11
- path database field, 15
- pathlen database field, 14
- pathseg.dat file, 12
- pathseg.keys file, 12
- proc database field, 15

R

- REGULAR field state, 11
- regular file, 6
 - defined, 9
- remigrating a file
 - internal processes, 16
- remove menu option, 74, 89
 - offline files, 59
 - partial-state files, 62
- removing or modifying a file
 - internal processes, 16
- report example
 - detecting discrepancies, 34
- report menu option, 32
- run_audit.sh, 3

S

- screen size
 - dmaudit interactive use, 21
- scrolling, 23
- search menu option, 44, 53
- snapshot menu option, 30
- snapshots
 - batch mode, 35
 - interactive, 30
 - resource requirements, 29
 - taking, 19
- soft-deleted, 2
- soft-deleted database entry
 - daemon database, 12
- state changes
 - within BFID sets, 6
- state field, 11

T

- transfer menu option, 74

U

- ufile_action output, 93
- ufile_data output, 91
- ufile_error output, 93
- ufile_name output, 92
- UNMIGRATING field state, 11
- unrecoverable errors
 - offline files, 56

- partial-state files, 60
- up menu option, 30
- user files
 - fixing, 50
 - fixing BFID sets without, 45
 - same BFID and different size, 64
 - same BFID and size, 67
- userid database field, 14
- utime database field, 14

V

- verifymsp menu option, 33
- view menu option, 25
- VISUAL environment variable
 - dmaudit interactive use, 21
- voided-state file
 - defined, 10

W

- workdir menu option, 84
- working directory
 - defined, 19

X

- xfsdump, 3
- xfsrestore, 3, 12