

IBM FileNet Image Services
Version 4.2

*Third Party Backup and Restore
Guidelines*



IBM FileNet Image Services
Version 4.2

*Third Party Backup and Restore
Guidelines*



Note

Before using this information and the product it supports, read the information in "Notices" on page 183.

This edition applies to version 4.2 of IBM FileNet Image Services (product number 5724-R95) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 1984, 2011.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this manual 12

Overview 12

Document revision history 14

Audience 14

Related documents 14

Accessing IBM FileNet documentation 14

Software education 15

Feedback 15

Documentation feedback 15

Product consumability feedback 15

1 Requirements and recommendations 16

Backup/restore methods 16

Simple backup/restore method 17

Advanced backup/restore method 17

Requirements for using a third-party product 18

Consistent backup/restore method 18

Performing offline backups and restores 19

 MKF databases 19

 RDBMS databases 20

 Mixing offline and online backups 20

Handling raw partitions or raw disks 20

| | |
|---|-----------|
| Handling capture settings (If FileNet Capture Professional is used) | 20 |
| Handling links or character special files | 21 |
| FileNet dataset synchronization | 21 |
| Customer support issues | 22 |
| Testing third-party backup/restore | 22 |
| Disk mirroring | 24 |
| Backup frequency | 25 |
| Enterprise backup/restore | 26 |

2 Simple Backup and Restore 27

| | |
|---|----|
| Simple restore summary | 28 |
| Recovering lost documents after restoring | 30 |
| Simple backup procedures | 31 |
| Simple restore procedures | 33 |
| Multiserver configuration | 43 |

3 Advanced backup methods and recommendations 44

| | |
|--|-----------|
| Backing up FileNet Image Services software and datasets | 45 |
| Summary backup table | 45 |
| File system software and data | 47 |
| Detailed backup table | 50 |
| Databases and cache | 52 |
| Permanent database | 53 |
| Permanent database recovery log | 53 |

| | |
|--|----|
| Security database | 54 |
| Security database recovery log | 54 |
| Transient database | 55 |
| Transient database recovery log | 55 |
| Cache | 56 |
| Index database | 57 |
| UNIX Operating Environment | 58 |
| FileNet Image Services for Windows Server | 58 |
| Index database redo log or transaction log | 60 |
| Archived logs | 60 |
| FileNet Image Services for UNIX platforms | 60 |
| FileNet Image Services for Windows Server | 61 |
| WorkFlo Queue database | 62 |

Backing up Microsoft SQL Server databases 62

Backing up non-FileNet datasets 64

| | |
|---|----|
| FileNet Image Services for AIX/6000 – backup/restore tips | 65 |
| FileNet Image Services for Windows Server – backup/restore tips | 66 |

Links and character special files 67

| | |
|--|----|
| Examples of links and character special files | 68 |
| FileNet Image Services for AIX/6000 | 68 |
| FileNet Image Services for HP-UX | 68 |
| FileNet Image Services for Solaris | 69 |
| Backing up and restoring /fnsw/dev/1 | 71 |
| Symbolic links | 71 |
| Character special files | 72 |
| No links | 72 |
| Raw partition and /fnsw/dev/1 names by server platform | 73 |
| FileNet Image Services for AIX/6000 | 73 |
| FileNet Image Services for HP-UX | 74 |

| | |
|-------------------------------------|-----------|
| FileNet Image Services for Solaris | 75 |
| Finding raw partition names | 76 |
| FileNet Image Services for AIX/6000 | 76 |
| FileNet Image Services for HP-UX | 77 |
| FileNet Image Services for Solaris | 78 |
| Advanced backup methods | 80 |

4 Advanced restore methods 83

| | |
|---|------------|
| Summary of advanced restore task order | 84 |
| Single dataset restore task order | 85 |
| Multiple dataset restore task order | 85 |
| Multiple dataset restore examples | 85 |
| Requirements for a restore | 88 |
| Restore methods by dataset type | 89 |
| Permanent database | 90 |
| Permanent database recovery log | 99 |
| Security database | 103 |
| Security database recovery log | 105 |
| Index database | 109 |
| RDBMS redo/transaction log | 113 |
| WorkFlo Queue database | 114 |
| WorkFlo Queue database recovery log | 118 |
| Transient database and cache – restore | 122 |
| Transient database – initialize | 127 |
| Transient database recovery log | 131 |
| Multiserver restore | 136 |
| Stopping and restarting each server | 136 |

Procedure for a Visual WorkFlo database 137

Restore issues 138

5 Synchronization and recovery information 139

FileNet dataset synchronization 140

FileNet dataset sequence 140

Synchronization requirements and examples 141

Synchronization problems 144

Lost data 144

Duplicate document numbers 144

Background information on ID numbers 144

snt.chkpt file (SNT checkpoint) 145

Updating the scalar numbers table 146

Running SNT_update 147

Updating the snt.chkpt file 149

Importing documents from optical media 149

Import issues 150

Import assumptions 151

Import diagram 151

MKF database rollforward recovery 154

MKF rollforward process 154

FileNet MKF coordination 155

MKF summary 156

MKF recovery log requirements 158

Independent software vendor product coordination with MKF 158

Synchronization and recovery issues 160

Permanent database issues 160

Transient database and cache issues 160

Reinitializing transient database and cache 161

Restoring transient database 162

Restoring disk cache 162

Index database recovery issues 162

FileNet systems without optical media 163

Appendix A – Support policy for backup/restore 164

Precautions 164

Who to contact 165

Recommendation 165

Appendix B – FileNet system tools 166

deldocs 167

killfnsw 171

MKF_zeroaij 174

odrecover 176

Notices 183

Trademarks 187

U.S. Patents Disclosure 187

Index 188

About this manual

Third-Party Backup/Restore Guidelines presents recommendations and procedural guidelines for backing up and restoring FileNet® Image Services systems with third-party products.

This manual also describes FileNet dataset synchronization issues to consider when using a non-FileNet backup and restore software package.

Information specific to third-party products is not described in this manual. Such information must be obtained from the vendors.

Overview

This manual contains the following chapters and appendixes:

Chapter 1. Recommendations and Requirements

Explains criteria for choosing one of two third-party backup methods: Simple Backup/Restore or an Advanced Backup/Restore. Describes customer support issues. Recommends testing of the third-party product, disk mirroring, backup frequency, and FileNet Enterprise Backup and Restore. Explains requirements of using third-party software: offline backups and restores, capability of software to back up either raw disk or raw partitions, and capability of software to restore links and character special files.

Chapter 2. Simple Backup and Restore Method

Describes procedures for performing a raw disk backup and restore.

Chapter 3. Advanced Backup Recommendations and Procedures

Presents the recommended FileNet and non-FileNet datasets to back up. Describes backing up through symbolic links and character special files. Explains advanced backup procedures.

Chapter 4. Advanced Restore Methods

Describes advanced restore methods for performing full and partial restores by individual dataset using third-party software capable of restoring raw partitions. Presents restore procedures by type of dataset lost.

Chapter 5. Synchronization and Recovery Information

Discusses FileNet dataset synchronization issues, Multi-Key File (MKF) rollforward recovery issues, MKF coordination with backup/restore software, import concerns, using the SNT_update program, and guidelines for systems without optical storage.

Appendix A. Support for Backup/Restore

Explains support considerations for customers who use third-party backup and restore software.

Appendix B. FileNet System Tools

Describes the FileNet system tools used in backup and restore operations.

Document revision history

| Image Services version | Date | Comment |
|------------------------|----------|------------------|
| 4.2 | May 2011 | Initial release. |

Audience

This manual is intended for customers who are using non-FileNet backup and restore software. This document assumes that the operator has knowledge and experience in:

- UNIX® or Windows® Server system administration.
- FileNet Image Services system administration or support.
- Volume management.
- The specific third-party backup software product deployed by the operator.
- Oracle® Database or Microsoft® SQL Server administration.

Related documents

In addition to reading this manual, familiarize yourself with the documentation for your backup and restore software.

Accessing IBM FileNet documentation

To access documentation for IBM FileNet Image Services products:

- 1 On the www.ibm.com website, enter “Image Services Documentation” in the search box on the menu bar.

- 2 Select **IBM - Product Documentation for FileNet Image Services** from the list of search results.

Software education

IBM provides various forms of education. Please visit the IBM Training and Certification for IBM software page at:

www.ibm.com/software/sw-training.

Feedback

Documentation feedback

Send comments on this publication or other IBM FileNet Image Services documentation by e-mail to comments@us.ibm.com. Be sure to include the name of the product, the version number of the product, and the name and part number of the book (if applicable). If you are commenting on specific text, include the location of the text (for example, a help topic title, a chapter and section title, a table number, or a page number).

Product consumability feedback

Help us identify product enhancements by taking a [Consumability Survey](#). The results of this comprehensive survey are used by product development teams when planning future releases. Although we are especially interested in survey responses regarding the most recent product releases, we welcome your feedback on any of our products.

The survey takes approximately 30 minutes to complete and must be completed in a single session; there is no option to save a partially completed response.

Requirements and recommendations

This chapter provides information you need to know before you use a third-party backup product. Read the [“Backup/restore methods” on page 16](#) to understand the criteria for choosing a Simple Backup/Restore or an Advanced Backup/Restore.

Also read the recommendations and requirements in this chapter regarding:

- Requirements when using a third-party product
- Customer support issues
- Testing the third-party product
- Disk mirroring
- Backup frequency
- Enterprise Backup/Restore (EBR)

Backup/restore methods

There are two backup/restore methods: Simple and Advanced. You should choose a backup/restore method based on the following criteria:

- The capabilities of your third-party software. Will it do full disk backups or partial (dataset specific) backups?
- Whether or not you wish to back up and restore individual FileNet datasets.

Simple backup/restore method

The Simple method backs up and restores the entire disk(s) which contain the datasets. It is the easiest and recommended method. Procedures for the Simple method are found in [**“Simple Backup and Restore” on page 27.**](#)

Advanced backup/restore method

The Advanced method allows you to specify individual datasets or files. This method is more complex in restore scenarios because more activities have to be considered when restoring individual datasets. Some of the considerations are:

- Full restore of individual datasets and rescan/re-entry of lost documents.
- Multiple or single restore of individual datasets and import from optical media to recover lost documents.
- Rollforward of the index database if you only need to restore the index database.
- Reinitialization of the transient database if you only need to recover data in the transient database and cache.

Procedures for the Advanced method are found in [**“Advanced backup methods and recommendations” on page 44.**](#)

If your third-party backup software cannot back up individual datasets or files then you must use the Simple method.

Requirements for using a third-party product

If your policy is to use a third-party backup and restore package rather than the FileNet EBR program, consider the following requirements:

- You must be consistent with your backup/restore method. You cannot “mix” Simple and Advanced tapes.
- You must perform offline backups and restores.
- Your third-party software must be able to back up and restore raw partitions or an entire disk.

If your third-party software does not have this capability, your software is not qualified to back up FileNet Image Services software and databases. The exception is that FileNet Image Services for Windows Server backs up and restores data in files, not raw partitions.

- Your third-party software should restore links or character special files during a restore if necessary (UNIX systems only).
- You must address database synchronization issues.

Consistent backup/restore method

You cannot “mix” Simple and Advanced tapes. If you use the Simple backup method (entire disk), you must do an entire disk restore. Entire disk (Simple) Backup tapes cannot be restored selectively by dataset. The Advanced Restore method allows you to restore specific datasets provided you did your backup with the Advanced Restore method.

Performing offline backups and restores

You are required to perform third party backups and restores **offline** due to special characteristics of the MKF databases and MKF recovery logs.

MKF databases

Third-party backup software does not have the capability to interact and coordinate with FileNet MKF database operations in the manner of FileNet EBR backup program.

The MKF database operations use MKF recovery logs to track all MKF database changes since any of the two following events occur:

- The last backup made with EBR.
- The last time FileNet Image Services software was brought down.

The MKF recovery logs write (or “flush”) the changes to the MKF databases when FileNet Image Services is brought down or when a backup is made using FileNet Image Services EBR backup program.

Using an **online** third-party backup/restore means that the recovery logs may rarely write to the MKF databases. The recovery logs may outgrow and overwrite themselves thus wiping out all pending MKF changes. This means MKF recovery logs may fall out of synchronization with the MKF databases in a third-party backup scenario. This will cause difficulties when you restore because FileNet Image Services will automatically attempt to apply the MKF recovery log to the MKF databases when FileNet Image Services is started after a restoration.

RDBMS databases

Although it is possible back up RDBMS databases online with third-party packages that support online backups, this is not supported or recommended.

Mixing offline and online backups

Do not mix an offline backup with an online backup due to database synchronization requirements. For example, do not restore an offline backup of the MKF databases and an online backup of the Oracle index database together.

Handling raw partitions or raw disks

If you need to back up and restore individual datasets, your third-party software must have the capability to back up and restore specific raw data partitions because UNIX-based FileNet Image Services datasets are located in raw partitions.

Important

FileNet Image Services for Windows Server stores data in files.

If you do not need to back up and restore individual datasets, your third-party software should have the capability to back up and restore an entire raw disk. You need to back up and restore all raw disks containing all FileNet datasets.

Handling capture settings (If FileNet Capture Professional is used)

Capture settings normally reside in cache. In the event of a corrupted transient database or cache, you may lose your Capture settings. To ensure a good copy of the Capture settings, you should back them up

to a local repository as part of a regular backup schedule. See your FileNet Capture Professional documentation for more information on these procedures.

Handling links or character special files

UNIX-based FileNet systems are set up to use symbolic links or character special files. A link is a character special file that points to data stored in a different physical location. If your FileNet system uses symbolic links or character special files, your third-party software must be able to restore those links or character special files during a restore. If your software lacks that capability, you must recreate the links or character special files in some other way, such as utilizing platform-specific command line utilities. See [**“Links and character special files” on page 67**](#) for more information.

FileNet dataset synchronization

The FileNet system employs several datasets to track documents. These datasets must be synchronized, that is, updated to the same point in time. Non-synchronization will cause data integrity problems or put your FileNet Image Services software in a non-operational state. Your FileNet datasets become vulnerable to these synchronization hazards during the restore process.

Specifically, lost data and duplicate document numbers will result from unsynchronized datasets. Your goals after a restore are to synchronize all MKF and non-MKF datasets and recover/re-enter any data lost on your FileNet system. Refer to [**Chapter 5, “Synchronization and recovery information,” on page 139**](#) for more details.

Customer support issues

You are responsible for the protection and integrity of your data when you use a third-party product to back up and restore your FileNet datasets. You need to contact your third-party backup product vendor on support issues.

Read about support for third-party backup and restore in [**“Appendix A – Support policy for backup/restore” on page 164**](#) of this manual. If you are not comfortable with the support conditions, it is recommended you use EBR.

Testing third-party backup/restore

To ensure a usable backup, you should test your third-party product on a test system **before** implementing the third-party product on your production system. In addition, you should make a backup of your system every six months with FileNet EBR. This will ensure that the foundation of your databases can be rebuilt from FileNet backups, if necessary. Maintaining reasonably current FileNet backups could reduce the length of recovery time of your databases. Store your FileNet backups either at an off-site location or in a fireproof safe.

Follow these steps to test your third-party product:

- 1 Perform a complete backup using FileNet EBR.
- 2 Perform a backup using your third-party product.
- 3 Scan and commit new documents after the backup.
- 4 Corrupt or delete one or more datasets.

For example, remove the permanent database logical volume (fn_permanent_db0), using the appropriate operating system tool.

Attention

Do not use a FileNet system tool to delete any documents as part of this test because the deletions are recorded in the MKF recovery log or RDBMS redo log. These deletions will be automatically reapplied from the recovery or redo log to a restored database (unless you zero out the contents of the logs before the restore).

Instead of deleting documents from a database, corrupt the data in a database using an appropriate operating system tool, for example, dd for UNIX platforms. Corrupting the data prevents the database from starting properly.

- 5 Start FileNet Image Services software using the Task Manager Start button or by using the command line.

During FileNet Image Services startup, verify that you receive fatal error messages (similar to a disk crash). This confirms that you successfully corrupted or destroyed datasets.

- 6 Stop the FileNet Image Services software using the Task Manager Stop button.
- 7 Restore from backup tapes created by the third-party product and recover any lost data.

To recover lost data, follow the Simple Restore or Advanced Restore procedures in [Chapter 2, “Simple Backup and Restore,” on page 27](#) or [Chapter 4, “Advanced restore methods,” on page 83](#).

- 8 Start your FileNet system using the Task Manager Start button or by using the command line.

If no errors occur during startup, your restore was successful.

- 9 Check your system to ensure that all data and documents were recovered and synchronized.

Check that the deleted datasets were recovered from backup tapes.

Disk mirroring

Your system should be set up with **disk mirroring** (also called multiple disk copies) or a **RAID** (Redundant Array of Independent Disks) sub-system. All your magnetic disks should have at least one mirrored copy. Disk mirroring protects uptime by maintaining system availability in the event of a disk failure. If a disk fails, disk mirroring prevents the loss of work performed since the last backup and allows your system to continue processing. Mirroring is also the most feasible way to **protect cache**.

Normally only one disk spindle crashes. Unless both disk spindles of a mirrored pair crash, the system continues to operate without interruption (but it does depend on which disk is the disk that crashes).

Mirroring does not negatively impact overall system performance. The performance of reads, which usually occur more frequently than writes, is actually increased and the performance of writes is only slightly reduced.

Mirroring is not a substitute for normal backups. Mirroring and backups complement each other. A site disaster, such as an earthquake, flood, fire, or sabotage, could destroy both members of a mirrored disk pair. Backup tapes kept physically separate from the computer system provide additional protection against disasters.

Backup frequency

The timing and frequency of backups is site-dependent and is usually determined by the number of daily transactions and the business requirements for these transactions. If you are using third-party backup software, the recommended backup frequency of the FileNet databases is **daily full backups**. Daily full backups ensure that you do not lose more than one day's worth of work in the event of a system failure.

The FileNet Image Services software itself only needs monthly backups because this software does not change frequently.

See [**“Backing up FileNet Image Services software and datasets” on page 45**](#) for backup recommendations for the FileNet file system release software, databases, and cache.

A full system backup is recommended immediately before and after the following events:

- Reconfiguration
- Software installation
- Operating system updates
- FileNet Image Services software updates
- Running MKF_fixup or MKF_fixone
- Dropping or adding indexes in relational databases

Note

MKF_fixup and MKF_fixone should only be run by FileNet support personnel.

Enterprise backup/restore

If you are investigating alternative backup and restore software packages, consider FileNet Enterprise Backup/Restore (EBR) product.

EBR is designed to keep datasets synchronized and prevent loss of data after a restore. EBR performs online backups and automatically rolls forward all databases to the end of the last transaction whenever possible. For this reason, EBR may be a better solution for 7 by 24 operations (7 days a week, 24 hours a day), where the system cannot be stopped for backups.

EBR is capable of:

- Enterprise-wide offline or online backups and restores of multiple FileNet systems and servers
- Striped dataset backups
- Data compression
- Pipelining
- Full and interval database backups
- Backing up to a disk file or to a tape device

Simple Backup and Restore

This chapter describes the Simple Backup/Restore method. The simple method is only for third-party software packages that have the capability to back up and restore entire disks.

Use the Simple Backup/Restore method if both of the following items are true:

- The third-party software is capable of backing up and restoring entire raw disks.
- You do not need to back up and restore individual datasets.

When backing up an entire disk(s), you do not need detailed information about database and cache dataset types, file system software, individual dataset and raw partition names, and link and character device file names. The restore process should restore any links or character device files used in the FileNet system. Restoration should not have synchronization problems with your magnetic disk datasets (databases and cache).

For a backup and restore of individual datasets and for a description of more advanced methods, read [Chapter 3, “Advanced backup methods and recommendations,” on page 44](#) and [Chapter 4, “Advanced restore methods,” on page 83](#).

Important

To ensure that you can restore your backups, it is imperative that you perform a **practice backup and restore on a test system**. Verify that

your restore software works, especially for MKF and index database restores. See [“Testing third-party backup/restore” on page 22](#) for more information.

As a safety net, you should make a full backup of your system with FileNet EBR software **before** using the third-party backup/restore package. If the third-party software fails to restore, you can restore with the FileNet backup and rebuild the foundation of your databases. Store your FileNet backups either at an off-site location or in a fireproof safe.

Simple restore summary

The Simple Restore approach for restoring your datasets is described below.

Simple Restore Summary

| Factor | Explanation |
|----------------------------------|--|
| Assumptions | Raw disk backups of all datasets are available. |
| Datasets restored | Full restore of MKF databases (permanent, transient, security), index database, cache, and file systems. |
| Reset of the SNT before scanning | Advance the Scalar Numbers Table (SNT) after the restore through the use of SNT_update. |
| Lost document recovery method | Rescan, reindex, recommit all documents not on the restore tape. Import from optical, if available. |

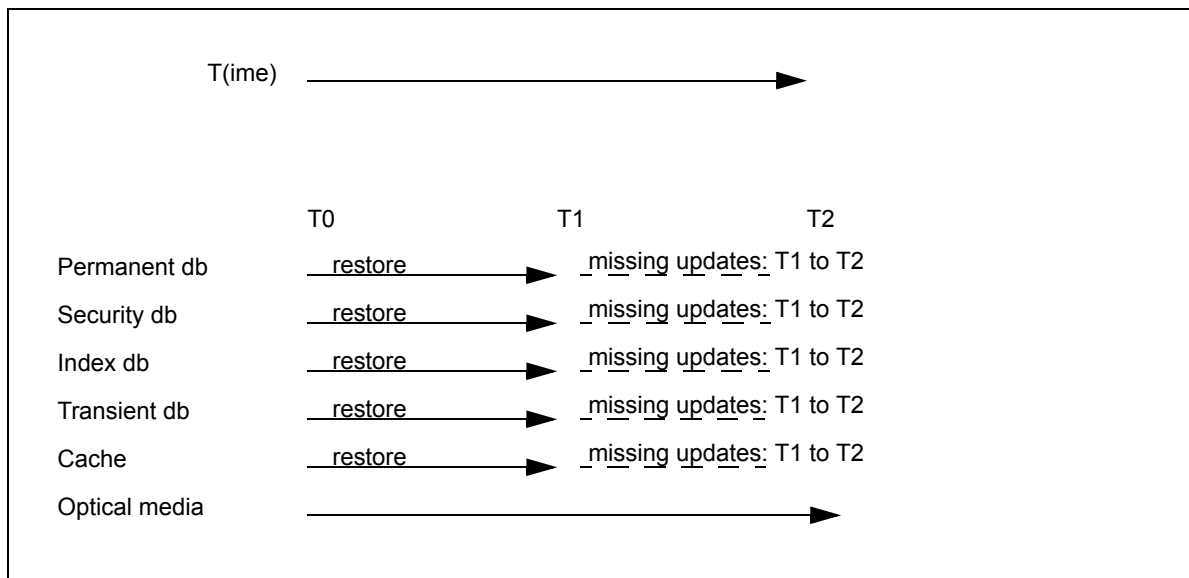
Simple Restore Summary, Continued

| Factor | Explanation |
|---------------|---|
| Advantages | Databases are synchronized with each other. |
| Disadvantages | Although the databases and cache are synchronized with each other after the restore, all the magnetic disk datasets and cache may be out of synchronization with the optical media, i.e., documents entered after the last backup may exist on optical media but referenced in the permanent and index databases. MKF recovery logs will not rollforward because they have been restored. You need to rescan all documents not on the restore tape. |

Note If you are using third-party restore software, you are required to perform all backups and restores **offline**.

The Simple Restore is graphically represented in the following figure. The figure assumes that a raw disk backup was performed at time T1. The figure portrays a system failure at time T2 and a full raw disk restore of the magnetic disk datasets up to time T1. The missing

updates from time T1 to T2 need to be rescanned (or imported from optical media) to be up-to-date with data on optical media.



Simple Restore: Raw Disk Restore and Rescan

Recovering lost documents after restoring

After the restore, there will likely be lost documents, that is documents not found on the backup tape. This is due to documents that have been entered since your last backup was performed. There are two methods you can use to recover these lost documents:

- Rescanning, reindexing, and recommitting all documents entered since the last backup.
- Import the documents from optical media, as described in **[“Importing documents from optical media” on page 149](#)**. This

process reestablishes the permanent and index database records of those documents. The import procedure requires that you rescan any lost documents that were not migrated to optical disk, such as locked documents in page cache.

Simple backup procedures

You should perform full, daily backups. You must do backups **offline**. Online backups can cause MKF databases corruption when restored.

Important

If you use the Simple Backup method, you must use Simple Restore. You cannot selectively restore datasets when you have done a backup of the entire disk.

- 1 Stop the FileNet Image Services software by doing one of the following.

Tip

You must be a member of the fnadmin or fnop and fnusr groups to stop/start Image Services software.

- On UNIX-based Image Services platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line.

For FileNet-controlled databases, always use the Task Manager or command line to stop both FileNet Image Services software and the RDBMS database. Do **not** use the Microsoft SQL Service Manager to stop the RDBMS. If you do, an error displays because the Microsoft SQL Service is dependent on the Image Services ControlService.

Tip If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

If normal attempts fail to stop the system, evaluate your software state. Be sure that it is not that your system is running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [“killfnsw” on page 171](#), then run the `initfnsw stop` command.

Important Rebooting your operating system may automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the `initfnsw stop` command or the Stop button in Task Manager before you back up or restore.

- 2 Perform a backup of all your disks to tape using your third-party backup software.
- 3 After the backups are completed, start the Image Services software by doing one of the following:

Tip If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On UNIX-based Image Services platforms, enter:

Xtaskman &

- Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)
- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfns -y start** at the command line.

Simple restore procedures

A full disk restore recovers everything that was on the disk: operating system, file systems, logical volumes, links or device files, and RDBMS software.

Important

If you use Simple Backup, you must use Simple Restore. You cannot selectively restore datasets from a full disk backup.

- 1 Stop the FileNet Image Services software by using one of the following methods.

Tip

You must be a member of the fnadmin or fnop and fnusr groups to stop/start the Image Services software.

- On UNIX-based Image Services platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line.

For FileNet-controlled databases, always use the Task Manager or command line to stop both FileNet Image Services software and the RDBMS database. Do **not** use the Microsoft SQL Service Manager to stop the RDBMS. If you do, an error displays because the Microsoft SQL Service is dependent on the Image Services ControlService.

Tip

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

If normal attempts fail to stop the system, evaluate your software state. Be sure that it is not that your system is running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [**“killfnsw” on page 171**](#), then run the **initfnsw stop** command.

Important

Rebooting your operating system may automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always shutdown FileNet Image Services software with the **initfnsw stop** command or the Stop button in Task Manager before you back up or restore.

- 2 Write down the document ID number, optical media surface ID number, and background request number located in the scalar numbers table checkpoint file, if possible (the disk on which this file resides must be intact).

Tip If you cannot perform step 2 because the magnetic disk on which the `snt.chkpt` file is located has crashed, you need to perform step 6.

- a Change to the `/fnsw/local/sd` directory (or `\fnsw_loc\sd` directory for Image Services for Windows Server).
- b View the `snt.chkpt` file. For UNIX systems, enter the following on the command line:

cat snt.chkpt

For Windows Server, enter the following on the command line:

type snt.chkpt

The screen output resembles the following (your screen output will not have notes indicating what the numbers mean):

| | | | | | |
|-----------|-------------------|-----|---------------------------|-------|------------|
| 920030 | 4020 | 120 | 856828856 | _____ | Time Stamp |
| | | | Background Request Number | | |
| Document | Optical Media | | | | |
| ID Number | Surface ID Number | | | | |

The fourth number represents the date and time of the last update to the snt.chkpt file and translates into a time stamp when you run the SNT_update program.

- c Write down the first three numbers.

You need to note these three numbers before the restore because the restore may overwrite the current, valid checkpoint file. After the restore, you need to advance these numbers in the scalar numbers table to avoid generating any potential duplicate document ID, surface ID, and background request numbers when scanning resumes. For more information, see [**“Updating the scalar numbers table” on page 146.**](#)

- 3 Perform an entire disk restore of all disks from tapes made with your third-party backup software.
- 4 Reboot your system. Make sure that the operating system is fully functional.
- 5 Edit the scalar numbers table snt.chkpt file. This will advance the document ID, optical media surface ID, and background request numbers.
 - a Change to the /fnsw/local/sd directory (or \fnsw_loc\sd directory on Image Services for Windows Server).
 - b Modify the snt.chkpt file using an appropriate platform-specific text editor. On UNIX platforms, enter:

vi snt.chkpt

On Image Services for Windows Server, enter:

edit snt.chkpt

Modify the document ID number, optical media surface ID number, and background request number, using the exact numbers you wrote down in step 2.

- c Save the file.
- d Skip to **Step 7 on page 38**

Important

Advancing the document ID number in the scalar numbers table is critical to maintaining data integrity. This will prevent the system from creating duplicate document ID numbers in the permanent database. Duplicate document ID numbers can lead to unpredictable results.

- 6 Perform this step only if you were unable to perform step 2. This assumes the magnetic disk on which the snt.chkpt file is located has crashed. Otherwise, skip to **Step 7 on page 38**.
 - a You must estimate how much to advance the document ID number in the snt.chkpt file.

For example, if you typically scan 20,000 documents per day and your backup is two days old, you should estimate that your document ID number would increase by two days' worth of documents. Then add one extra day's worth of document IDs to this number (for example, $40,000 + 20,000$) to ensure that you do not underestimate. This means you should advance the document ID number by 60,000 documents. It is better to increase by too much than by too little.
 - b Edit snt.chkpt with your estimated document ID number and save the file.

Procedures for advancing the optical media surface ID number are described in [Step 11 on page 41](#) and [Step 12 on page 41](#).

You may leave the old background request number as is because duplicate background request numbers do not matter since the numbers are temporary.

- 7 Start the FileNet databases (permanent, transient, security, and index databases) using the `fn_util startdb` Image Services command. You must be a member of the `fnadmin` or `fnop` and `fnusr` and `dba` groups to do this.

Run the command from the system console on each Image Services server. Enter the following at the server command line:

`fn_util startdb`

Although the databases are started, FileNet Image Services software is not. This is to prevent users from entering documents at this recovery phase.

- 8 Run `SNT_update` immediately, using the following procedures:

- a Log on as the `fnsu` user.
- b Enter the command:

`SNT_update`

- c Enter **Y** when you are prompted as follows:

Sure you want to update the scalar numbers table? Y or N.

`SNT_update` displays messages indicating either successful completion or errors that prevented successful update of the database. For a

detailed description of the SNT_update program, see [**“Updating the scalar numbers table” on page 146.**](#)

- 9 Copy or rename the restored CDB file into the next highest numbered CDB file. Update the CDB file to reflect the system changes you entered since the last backup (instructions will follow).

Perform this step if you have changed your FileNet configuration database (CDB) file since the last backup, especially if the change was significant. This will ensure that your system state and the current CDB file are consistent.

Significant changes include adding new datasets or partitions. Other changes include adding system devices such as optical devices, printers, tape drives and changing tuning parameters such as memory buffers and disk space allocations

Tip Your configuration database file (CDB) changes when you run **fn_edit** (or Configuration Editor on NT) and save your changes.

The reason this step is included is that the restored system state may not match the current CDB file used by the Configuration Editor. The CDB file number is advanced to the next highest number each time you update your FileNet system configuration through the **fn_edit** utility. (Older files are not automatically deleted.) When you restart your system after a full restore, the Configuration Editor uses the highest numbered CDB file found in the directory.

For example, at the time you performed the last full backup, the CDB file number was 10 (IMS_10.cdb). Later you ran **fn_edit** to add extents and datasets. This action created and the current CDB file: IMS_11.cdb. After a full restore, the CDB file that is currently recognized by the system is IMS_10.cdb **not** IMS_11.cdb. IMS_11.cdb was overwritten in the restore and therefore, IMS_10.cdb has the highest

number. You should copy IMS_10.cdb into the next highest numbered file, IMS_12.cdb.

The CDB file is found in:

UNIX

/fnsw/local/sd/conf_db/IMS_nnn.cdb

WIN

<drive>:\fnsw_loc\sd\conf_db\IMS_nnn.cdb

- a Copy the CDB file by using the UNIX **cp** command or the Windows Server **copy** command (or File Windows Server Explorer method).
- b Use the fn_edit utility (or Configuration Editor on Windows Server) and re-enter changes made to the configuration since the last backup.
- c Synchronize the configuration files and the datasets by entering the following on the command line:

fn_build -a

- 10 Restart FileNet Image Services software using the method appropriate for your platform. Restart will stop and start up FileNet Image Services software.

Tip

Start the RDBMS database first, if you have a site-controlled RDBMS database, **before** starting FileNet Image Services software.

UNIX

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Restart** button from Task Manager. (The ampersand runs the process in the background.)

WIN

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Restart** button.

- 11 Update the optical media surface ID numbers in the checkpoint.osa data file. This file is located in the /fnsd/local/sd (or \fnsd_loc\sd) directory. The way you update this file is by running the **Identify Media in Library** option from the **Storage Library Control Program**. For more information about Storage Library Control, see the *System Administrator's Handbook*. To download IBM FileNet documentation from the IBM support page, see [**"Accessing IBM FileNet documentation" on page 14.**](#)

The program locates all known optical media and automatically updates the optical media surface ID number in the checkpoint.osa file with the new numbers. This step provides the input for the next step, which updates the surface ID number in the scalar numbers table snt.chkpt file.

- 12 Manually update the optical media surface ID number in the snt.chkpt file using input from the **DOC_tool** program.

The surface ID number must be updated **before** scanning resumes to prevent duplicate surface ID numbers on the next new optical media.

- a At a command line or MS-DOS prompt, enter:

DOC_tool

- b At the DOC_tool prompt, enter the subcommand:

dev

The dev subcommand displays the contents and surface ID for the gripper and each drive and slot in the storage library from the checkpoint.osa file.

Respond as follows to the prompts:

Print status of empty slots? (y/n) : **n**

Print status of disabled slots? (y/n) : **y**

Your display will look similar to:

| | | |
|---------|----------------|-------------|
| Gripper | Contents=known | Surfid=4002 |
| Drive 0 | Contents=known | Surfid=4000 |
| Drive 1 | Contents=known | Surfid=3998 |
| Slot 21 | Contents=known | Surfid=3996 |
| Slot 28 | Contents=known | Surfid=3965 |

- c Find the highest numbered surface ID from your display. Compare this number to the surface ID number in the snt.chkpt file.
- d If the surface ID in the snt.chkpt file is not two numbers higher than the surface ID displayed in DOC_tool, edit the snt.chkpt file to increase the surface ID number by two.

The snt.chkpt file's surface ID must be two numbers higher than the surface ID displayed in DOC_tool. For example, if the surface ID number in the snt.chkpt file is 4000 and DOC_tool displays a surface ID number of 4002, then manually update the surface ID number in snt.chkpt to 4004.

- e Rerun the SNT_update program by entering:

SNT_update

f Enter **Y** when prompted to update the scalar numbers table.

13 Rescan any lost documents not on the backup tape.

Identify work done since the last backup and after the restore. This should include documents that were never written to optical media, such as lost batches and batches that have not fully migrated to optical media. Rescan, reindex, and recommit those documents.

Tip If you rescan lost data, do not also import the same data from optical media. Otherwise the same document will occur twice, although with different document IDs.

Multiserver configuration

If you have a multiserver configuration, you may want to view the information in [**“Multiserver restore” on page 136**](#) for:

- Procedures for starting and stopping each server
- Additional restore procedure when your Oracle database contains a Visual WorkFlo database

3

Advanced backup methods and recommendations

This chapter applies to customers who use independent software vendor software packages that are capable of backing up individual FileNet datasets located in raw partitions. These methods and recommendations are considered advanced because backing up and restoring individual datasets is more complex than backing up an entire disk.

This chapter presents background and detailed information for an advanced backup and subsequent restore. Topics include:

- Summary and detailed description of the FileNet file system software and data, databases, and cache to back up
- Backing up Microsoft® SQL Server™ databases
- Backing up non-FileNet datasets
- Links and character special files
- Backing up and restoring /fnsd/dev/1 (or \fnsd\dev\1)
- Raw partition and fnsd directory names by server platform
- Finding raw partition names
- Advanced backup procedures

Backing up FileNet Image Services software and datasets

Read this section to gain both a summary and a detailed description of the FileNet file system software and data, databases, and cache recommended for backup.

If you are using independent software vendor backup software, you should perform full, daily backups. You must do **offline** backups. Doing online backups can cause restored MKF databases to be corrupted. Although the Relational Database Management System (RDBMS) database may be backed up online, this method is not recommended, even if your independent software vendor software provides that capability.

Summary backup table

The following table presents a summary of the FileNet Image Services software, databases, and cache that you should back up. Note that this also includes the RDBMS software. Names of raw partitions in the

table are determined by your system configuration. See [“Databases and cache” on page 52](#) for more information on raw partition names.

Summary of FileNet Software and Datasets to Back Up

| Dataset/File System | Type | Backup Frequency |
|---|---|------------------|
| /fnsw or <drive>:\fnsw | File system (UNIX) or directory (Windows Server) with FileNet Image Services software | Monthly |
| /fnsw/dev/1 or <drive>:\fnsw\dev\1 | Actual datasets for FileNet Image Services for Windows Server Links to data for FileNet Image Services for UNIX platforms. | Daily |
| /fnsw/local or <drive>:\fnsw_loc | File system (UNIX) or directory (Windows Server) with configuration files | Daily |
| /fnsw/local/tmp or <drive>:\fnsw_loc\tmp | File system (UNIX) or directory (Windows Server) with configuration files | Daily |
| /fnsw/local/sd or <drive>:\fnsw_loc\sd | File system (UNIX) or directory (Windows Server) with configuration files | Daily |
| /fnsw/local/oracle/control0 and /fnsw/local/oracle/control1 or <drive>:\fnsw_loc\oracle\control0 and <drive>:\fnsw_loc\oracle\control1 | Oracle control files. File system in AIX®. Files on all other platforms. | Daily |

Summary of FileNet Software and Datasets to Back Up, Continued

| Dataset/File System | Type | Backup Frequency |
|---|--|------------------|
| /usr/ora/817 <drive>:\orant | Directory with Oracle software | Monthly |
| <drive>:\MSSQL7 | Directory with Microsoft SQL Server software | Monthly |
| Permanent database | Raw partition (UNIX only) | Daily |
| Permanent database recovery logs | Raw partition (UNIX only) | Daily |
| Transient database | Raw partition (UNIX only) | Daily |
| Transient database recovery logs | Raw partition (UNIX only) | Daily |
| Security database | Raw partition (UNIX only) | Daily |
| Security database recovery logs | Raw partition (UNIX only) | Daily |
| Cache | Raw partition (UNIX only) | Daily |
| Index database | Raw partition (UNIX only) | Daily |
| Index database redo or transaction logs | Raw partition or file (UNIX only) | Daily |
| WorkFlo queue database | Raw partition (UNIX only) | Daily |
| WorkFlo queue database redo or transaction logs | Raw partition (UNIX only) | Daily |

File system software and data

This section lists the major FileNet file system software and data that you should back up. Many FileNet datasets for UNIX platforms are located in raw partitions or raw logical volumes. Your independent software vendor backup software must be able to back up and restore raw logical volumes.

Important

On FileNet Image Services for Windows Server, all databases are in files, not in raw partitions.

The number and names of the raw logical volumes on your system are determined by how you configure your system. The raw logical volume names shown in this chapter are merely examples.

Use the FileNet Configuration Editor (fn_edit) to identify all partitions defined in your FileNet system. Refer to the *System Configuration Editor Overview*, or to the System Configuration Tools online help. To download IBM FileNet documentation from the IBM support page, see **[“Accessing IBM FileNet documentation” on page 14.](#)**

Back up FileNet Image Services software in various subdirectories of the /fnsw directory on a **monthly** basis. Back up **daily** all other file systems containing, for example, data, system logs, event logs, and control files.

See the diagram below for a graphical representation of some of the subdirectories found under the /fnsw directory on a UNIX platform.

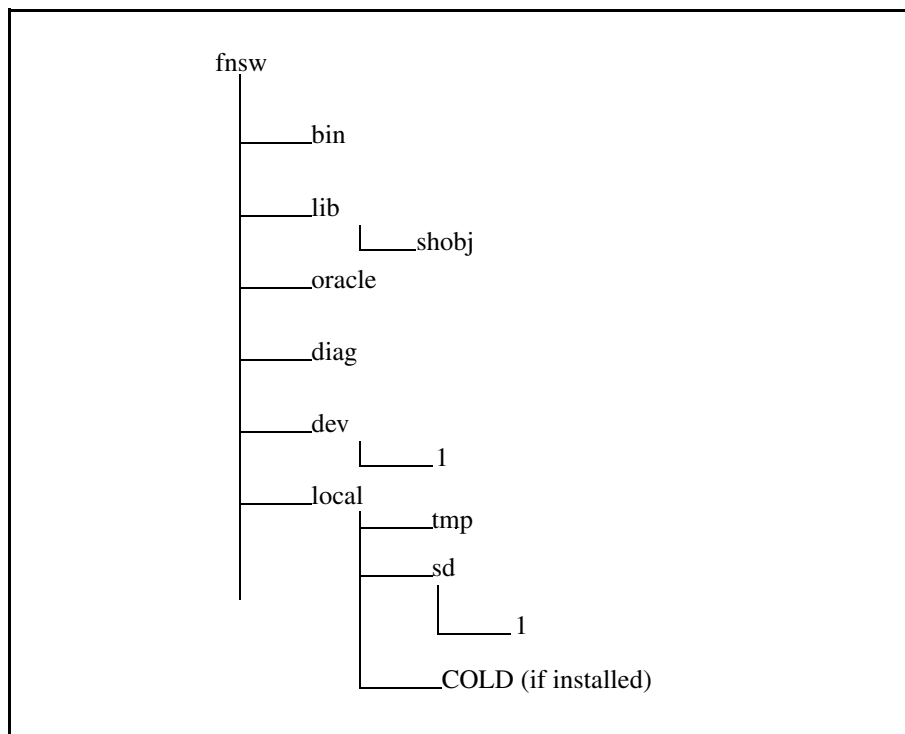


Diagram of File System Software

Detailed backup table

The table below describes the recommended file system software and data to be backed up and their backup frequency.

File System Software and Data to Back Up

| File System | Description | Backup Frequency |
|--|---|------------------|
| For UNIX-based systems, /fnsw including: /fnsw/bin /fnsw/lib/shobj /fnsw/oracle /fnsw/diag | The FileNet Image Services software is located under various subdirectories in the general /fnsw directory. | Monthly |
| On FileNet Image Services for Windows Server: <drive>:\fnsw | The FileNet Image Services software resides in a user-selected directory tree with a default name of \fnsw. It can be located on any logical drive and its name can be modified during FileNet Image Services installation. | |
| For UNIX-based systems: /fnsw/dev/1 | On all UNIX servers, the /fnsw/dev/1 directory contains pointers to raw partitions containing data in: <ul style="list-style-type: none">• MKF databases (permanent, transient, security)• RDBMS database• cache The actual data is in different physical locations and their directory names depend on your server platform. See <u>“Databases and cache” on page 52</u> for details. | Daily |

File System Software and Data to Back Up, Continued

| File System | Description | Backup Frequency |
|---|--|--|
| On FileNet Image Services for Windows Server: <drive>:\fnsw\dev\1 | The default location for MKF database, RDBMS, and cache files. The user determines on which disk drive to place MKF database and cache files. Use the FileNet utility, fn_edit, to identify file locations. | |
| /fnsw/local (for UNIX platforms) <drive>:\fnsw_loc (for Windows Server platforms) | Contains site-specific information, such as forms and server-dependent configuration files. This is a separate file system from /fnsw. The configuration database (CDB) is found in: /fnsw/local/sd/conf_db/IMS_nnn.cdb On Windows Server, the fnsw_loc files reside in a user-selected directory tree. The directory can be located on any logical drive and its name can be modified during FileNet Image Services installation. The configuration database (CDB) is found in: <drive>:\fnsw_loc\sd\conf_db\IMS_nnn.cdb | Daily and after changes to the configuration database. |
| /fnsw/local/tmp (for UNIX platforms) <drive>:\fnsw_loc\tmp (for Windows Server platforms) | Contains system logs, temporary files, performance monitoring files, and error logs. The \fnsw_loc\tmp files reside in a user-selected directory tree. | Daily |
| /fnsw/local/sd (for UNIX platforms) <drive>:\fnsw_loc\sd (for Windows Server platforms) | Contains the Network Clearinghouse (NCH) Database (NCH_db0). The \fnsw_loc\sd files reside in a user-selected directory tree. | Daily |
| /fnsw/local/oracle/control0 and /fnsw/local/oracle/control1 (for UNIX platforms) | Contain the Oracle control files for systems using an Oracle RDBMS. They must be backed up at the same time the index database is backed up. | Daily |

File System Software and Data to Back Up, Continued

| File System | Description | Backup Frequency |
|--|--|------------------|
| <drive>:\fnsw_local\oracle\control0 <drive>:\fnsw_local\oracle\control1 (for Windows Server platforms) | On Windows Server, the files reside in a user-selected directory tree. | |
| /usr/oracle (for pre-IDMIS 3.5.0 UNIX platforms) /usr/ora/xxxx (for IDMIS 3.5.0 and later UNIX platforms) | Contains Oracle software for systems using an Oracle RDBMS. | Monthly |
| <drive>:\orant (for Windows Server platforms) | On FileNet Image Services for Windows Server, the \usr\oracle files reside in a user-selected directory tree. | |
| <drive>:\MSSQL or <drive>:\MSSQL7 For MS SQL Server users | Contains Microsoft SQL Server software for systems using an SQL Server RDBMS (only on Windows Server). The file resides in a user-selected directory tree. | Monthly |
| Microsoft SQL Server databases | All databases in a Microsoft SQL Server RDBMS on FileNet Image Services for Windows Server. For a description, see <u>“Backing up Microsoft SQL Server databases” on page 62.</u> | Daily |

Databases and cache

This section describes the database and cache dataset types recommended for backup. Information specific to an operating environment is described as it affects the dataset. For information on the database and cache backup frequency, refer to the **“Summary of FileNet Software and Datasets to Back Up” on page 46.**

Important

The links to the actual data are always in the /fnsw/dev/1 directory on UNIX platforms. However the location of the actual data varies from

server to server. Just backing up /fnsw/dev/1 is not adequate because it contains only pointers and not the actual data. Your independent software vendor backup software must also back up the linked data. See [**“Links and character special files” on page 67**](#) for more information.

Permanent database

The MKF permanent database contains the mapping of document numbers to optical media locations as well as information about media families. The following is an example of a raw logical volume name on a UNIX platform:

```
/fnsw/dev/1/permanent_db0
```

Your configuration may include additional raw logical volumes, such as permanent_db1 and permanent_db2.

Important

Windows Server uses actual files and not raw logical volumes.

Permanent database recovery log

The MKF recovery logs associated with the permanent database contain a record of transactions to the database. As transactions occur, the data is written first to the recovery log and then to the permanent database.

Think of the recovery log as part of the database. Back it up at the same time you back up the database.

An MKF recovery log is important for database recovery. If the recovery log is available on disk, MKF database software applies the log automatically to the permanent database after the permanent database has been restored.

The following is an example of a raw logical volume name on a UNIX platform:

```
/fnsw/dev/1/permanent_r10
```

Security database

The security database contains FileNet security data.

The following is an example of a raw logical volume name on a UNIX platform:

```
/fnsw/dev/1/sec_db0
```

Security database recovery log

The MKF recovery log associated with the security database contains a record of transactions to the database. As the transactions occur, the data is written first to the recovery log and then to the security database. Think of the recovery log as part of the database and back it up at the same time you back up the database.

An MKF recovery log is important for database recovery. If the recovery log is available on disk, MKF database software applies the log automatically to the security database after the security database has been restored.

The following is an example of a raw logical volume name on a UNIX platform:

```
/fnsw/dev/1/sec_r10
```

Transient database

The transient database maps the contents of the magnetic disk cache and contains work-in-progress information for the FileNet system, such as batch status and read, write, and print requests.

The following is an example of a raw logical volume name on a UNIX platform:

```
/fnsw/dev/1/transient_db0
```

Your configuration may include additional raw logical volumes, such as transient_db1 and transient_db2.

Attention

If you back up cache, you must also back up the transient database because the transient database keeps track of all the objects in cache. Back up both the transient database and cache at the same time. Backing up cache without backing up the transient database would cause your cache to be out of synchronization with the transient database after a restore.

Transient database recovery log

The MKF recovery logs associated with the transient database contain a record of transactions to the database. As the transactions occur, the data is written first to the recovery log and then to the transient database. An MKF recovery log is important for database recovery. If the recovery log is available on disk, MKF database software applies the log automatically to the transient database after the transient database has been restored.

The following is an example of a raw logical volume name on a UNIX platform:

/fnsf/dev/1/transient_r10

Cache

Cache contains work in progress: images recently entered, documents recently retrieved, recently committed images that have not yet been written to optical media, uncommitted batches, images fetched to satisfy print requests, FileNet Capture settings, and FolderView data. On systems that never write to optical media, cache also contains committed documents.

The disk cache must be backed up if you store images permanently in cache instead of on optical media or have a significant number of uncommitted images in cache.

If you use FolderView cache and/or Revise cache, back up these caches because objects in these caches are not committed to optical media. Use the FileNet utility, CSM_exim. See the *FileNet Image Services System Administrator's Handbook*. To download IBM FileNet documentation from the IBM support page, see [**“Accessing IBM FileNet documentation” on page 14.**](#)

If you back up cache, you must also back up the transient database since the transient database keeps track of all the objects in cache. Backing up cache without backing up the transient database causes your cache to be out of synchronization with the transient database after a restore.

Important

You should also copy your Capture settings to a local repository. See your FileNet Capture Professional documentation for more information on these procedures.

On FileNet Image Services for AIX/6000, FileNet Image Services for HP-UX, and FileNet Image Services for Solaris, the disk cache is stored in raw logical volumes. Your backup software must be able to back up and restore raw logical volumes.

The following is an example of a raw logical volume name on a UNIX platform:

`/fnsw/dev/1/cachen`

where *n* is the cache logical volume number.

On FileNet Image Services for Windows Server, the disk cache is stored in files.

Index database

The description of the index database in this section is followed by platform-specific examples of raw logical volume names or file names.

The **index database** is created and managed by an RDBMS, for example, Oracle or Microsoft SQL Server. The index database keeps a record of every document in the system for searches and retrieval. The index database contains indexes, system-assigned document numbers, document class types, information that maps documents to folders, and the data dictionary.

If you do not use a dedicated server for WorkFlo queues, the index database also contains the WorkFlo queue tables. The index database and WorkFlo Queue database have identical raw data partition names.

UNIX Operating Environment

On FileNet Image Services for UNIX platforms, the index database is stored in raw logical volumes (or raw slices) under the /fnsf file system. Your backup software must be able to back up and restore raw logical volumes.

The following are examples of raw logical volume names for an Oracle database:

```
/fnsf/dev/1/oracle_db0  
/fnsf/dev/1/oracle_sys0 (if system tablespace is  
                        configured separately)
```

```
/fnsf/local/oracle/control0/ctl.ora  
/fnsf/local/oracle/control1/ctl.ora
```

```
/fnsf/dev/1/oracle_tr0  
/fnsf/dev/1/oracle_udb0 (full-use Oracle only)
```

Important

The oracle control files (ctl.ora) are actual (blocked) files and not raw character datasets.

FileNet Image Services for Windows Server

On FileNet Image Services for Windows Server, the index database is stored in an NT File System (NTFS) partition on your system. By default, the system configuration tools create the index database files under the \fnsf_loc directory tree. You can specify under which drive these directories will be located. Use the Data Sets window in the FileNet Image Services Configuration Editor to find the location of all FileNet Image Services datasets.

The following are examples of file names on FileNet Image Services for Windows Server with an Oracle database:

<drive>:\fnsw\dev\1\oracle_db0

<drive>:\fnsw\dev\1\oracle_sys0 (if system tablespace is configured separately)

<drive>:\fnsw_loc\oracle\control0\ctl.ora

<drive>:\fnsw_loc\oracle\control1\ctl.ora

<drive>:\fnsw\dev\1\oracle_tr0

<drive>:\fnsw\dev\1\oracle_udb0 (full-use Oracle only)

The following are examples of raw logical volume names for a Microsoft SQL Server 6.5 database:

\fnsw\dev\1\ms_db0.dat

\fnsw\dev\1\ms_db1.dat

\fnsw\dev\1\ms_tmp0.dat

\fnsw\dev\1\ms_udb0.dat (full-use or Visual WorkFlo only)

The following are examples of raw logical volume names for a Microsoft SQL Server 7 database:

\fnsw\dev\1\ms_db0.mdf

\fnsw\dev\1\ms_db01.mdf

Index database redo log or transaction log

When a transaction to an index database completes, the changes are written to a redo log (for an Oracle index database) or transaction log (for a Microsoft SQL Server index database).

For an Oracle database, only one redo log is used at a time and the redo log is assigned a sequence number for current data. When that log becomes full, the system switches to the next available log and increments the sequence number. For a Microsoft SQL Server database, the transaction log can optionally be dumped and truncated before it becomes full.

An Oracle index database has, by default, two redo logs (oracle_r10 and oracle_r11). A Microsoft SQL Server index database has one transaction log for each database (such as, ms_r10.dat or ms_r10.ldf for Microsoft SQL Server 7).

Think of the redo or transaction log as part of the database and back it up at the same time you back up the database.

Archived logs

Archived redo logs are an optional feature. An archived redo log or archived transaction log is a copy of the redo or transaction log. The archived redo log or archived transaction log is a separate file. It has been written to a specified location (backup device), such as magnetic disk or tape device. You should back up the archived redo or archived transaction logs as a part of your regular operating system backup.

FileNet Image Services for UNIX platforms

On FileNet Image Services for UNIX platforms, the index database redo or transaction log is stored in raw logical volumes (or raw slices)

under the /fnsf file system. Your backup software must be able to back up and restore raw logical volumes.

The following are examples of logical volume names for Oracle redo logs on UNIX platforms:

/fnsf/dev/1/oracle_rl0

/fnsf/dev/1/oracle_rl1

FileNet Image Services for Windows Server

On FileNet Image Services for Windows Server, the index database redo or transaction log is stored, by default, in files under \fnsf of the appropriate drive letter.

The following are examples of file names for Oracle redo logs:

<drive>:\fnsf\dev\1\oracle_rl0

<drive>:\fnsf\dev\1\oracle_rl1

The following are examples of file names for Microsoft SQL Server 6.5 transaction logs. The log is logically one transaction log, but the log can span two physical files.

<drive>:\fnsf\dev\1\ms_rl0.dat

<drive>:\fnsf\dev\1\ms_rl1.dat

<drive>:\fnsf\dev\1\ms_url0.dat (full-use or Visual
WorkFlo only file)

The following are examples of file names for Microsoft SQL Server 7 transaction logs.

```
<drive>:\fnsw\dev\1\ms_rl0.ldf  
<drive>:\fnsw\dev\1\ms_rl01.ldf
```

WorkFlo Queue database

The WorkFlo Queue database contains WorkFlo queue tables used by WorkFlo and managed by Oracle or Microsoft SQL Server. In smaller installations, the WorkFlo queue tables are usually stored in the index database. When a dedicated application server is used for WorkFlo queues and runs WorkFlo Queue services, the **WorkFlo Queue database** is located on the WorkFlo Queue server.

The index database and WorkFlo Queue database have identical raw logical volume names. See [“Index database” on page 57](#) for examples of raw logical volume names.

Backing up Microsoft SQL Server databases

The following table describes all the FileNet-created and SQL Server databases in a Microsoft SQL Server RDBMS and their backup frequency. Unlike Oracle, Microsoft SQL Server has many different databases residing in different locations. You should back up the SQL

Server databases listed in the physical devices column at the same time you do a full backup.

Microsoft SQL Server Databases to Back Up

| Database and Type | Logical Devices | Physical Devices | Contents | Frequency |
|--|---|---|--|---|
| fnsys MS SQL Server index database. FileNet-created. | fn_data <i>n</i> fn_log <i>n</i> (log) | ms_db0.dat ms_rl0.dat (log) | Indexing fields, document classes, document indexing status, folder information, SQL user data and WorkFlo queues. (If you have an Application server that contains SQL or WorkFlo Queue services, the SQL user data and WorkFlo queues reside in a separate database on the Application server.) Transaction log for fnsys database. | Daily The log is backed up with fnsys. |
| master MS SQL Server Master Database. Transaction log for master is located on the same physical/ logical device as master. SQL Server-created. | master | master.dat msdb.dat msdblog.dat (log) | System tables containing information on all databases, devices, SQL Server logins, space allocation, configuration, remote servers and logins, and server error-message language. Do not modify this database—see Caution on next page. Transaction log for the msdb database. | Daily. The log is backed up with the database. |

Microsoft SQL Server Databases to Back Up, Continued

| Database and Type | Logical Devices | Physical Devices | Contents | Frequency |
|--|---------------------------------|--------------------------------------|--|-----------|
| fnusr MS SQL Server User Database is for <i>full-use customers only</i> . FileNet-created. | usr_datan usr_logn (log) | ms_udb0.dat ms_url0.dat (log) | Tables containing user-specific data. Also contains System Tables common to each database. Transaction log for fnusr database | Daily |
| tempdb MS SQL Server Temporary Database; located on the master device. FileNet-modified. | tmp_datan | ms_tmpn.dat | Working storage and temporary tables. | Daily |
| model MS SQL Server Database; located on the master device | | | Defaults for creating user databases. | Daily |

CAUTION

Do **not** modify the Microsoft SQL Server Master database. The SQL Server Master database is configured automatically during installation. If you manually alter the Master database, you may not be able to successfully restore the database.

Backing up non-FileNet datasets

In addition to backing up and restoring FileNet datasets, you must also back up and restore non-FileNet datasets. Non-FileNet datasets are platform specific and include the vendor operating system, customer

applications, and other customer-specific software. The procedures for backing up and restoring non-FileNet datasets are described in documentation obtained from the platform vendors and are not described in this manual. Refer to your vendor documentation for more information.

FileNet installation and update procedures often involve modifying native operating system configuration files, especially in the /etc directory on FileNet Image Services on UNIX platforms. For information on operating system update procedures, refer to your platform vendor documentation.

FileNet-specific information may be added to the native server system configuration database, such as the Registry database on FileNet Image Services for Windows Server or the ODM database on FileNet Image Services for AIX/6000.

The following information regarding backing up and restoring on FileNet Image Services for AIX/6000 and FileNet Image Services for Windows Server is presented as a supplement to information found in your vendor documentation.

FileNet Image Services for AIX/6000 – backup/restore tips

On FileNet Image Services for AIX/6000 a FileNet system typically has two volume groups: the root volume group (rootvg), containing the FileNet file system, and the FileNet volume group (fnvg), containing all databases, Oracle software, and Oracle control files. However, some configurations have the Oracle control files in the root volume group. To see which logical volumes and file systems are in rootvg, enter the **lsvg -l rootvg** command at the console.

AIX/6000 users can use a mksysb tape to restore the root volume group. If your Oracle control files are in the root volume group, using a

mksysb tape to restore the root volume group also restores the Oracle control files. **This causes the Oracle control files to be out of synchronization with the Oracle index database, which is located in the FileNet volume group.** As a result, Oracle will not start.

If you use a mksysb tape to restore the root volume group, you can do one of the following to keep Oracle control files and the Oracle index database synchronized:

- Unmount the Oracle control files before making the mksysb tape.
When you restore using the mksysb tape, the Oracle control files will not be part of the restore.
- Restore the Oracle index database and its corresponding control files after restoring the root volume group.

FileNet Image Services for Windows Server – backup/restore tips

Windows Server servers have a Registry database called the NT Registry. If the Registry database is lost, use one of the following procedures to recover it:

- Boot up with a backup copy of the Registry:
During bootup, press the spacebar when the prompt “Press spacebar now to invoke last known good menu” displays. Windows Server boots up with a backup copy of the Registry.
- Reinstall Windows Server:
 - a Reinstall the Windows Server software.
 - b Import a backup of the Registry database.

- c Restore the rest of the system.
- Use Emergency Repair Disk.

Links and character special files

If your FileNet Image Services system is on a UNIX platform, it will have symbolic links or character special files. Links or character special files point to raw partition storage areas. In this case, the actual raw partition data are the FileNet Image Services databases: permanent, transient, and security, index database, and cache. FileNet Image Services software uses the links and character special files to back up and restore the actual data.

FileNet Image Services for AIX/6000 and FileNet Image Services for Solaris use symbolic links. FileNet Image Services for HP-UX uses character special files. Symbolic links and character special files reside in the `/fnsd/dev/1` directory. The raw partition names depend on your server platform. Because data is stored in raw partitions, backup software should be able to back up and restore raw data partitions if you need to back up and restore individual databases.

FileNet Image Services for Windows Server does not use links. The actual data is stored in files.

If your system uses links or character special files, it is not adequate to back up only the data. Your independent software vendor software must be able to recreate the links or character special files during a restore. If your software does not have that capability, the links or character special files must be recreated some other way, for example, by using platform-specific commands or utilities.

Examples of links and character special files

The following examples, by server platform, show how FileNet Image Services software looks for data stored in raw partitions through links and character special files. Your raw partition names may be different from the ones in the examples.

FileNet Image Services for AIX/6000

The `/fnsw/dev/1` directory contains symbolic links to the logical volumes in `/dev`.

For example:

`/fnsw/dev/1/transient_db0`

points to

`/dev/rfn_trans_db0`

where `/fnsw/dev/1/transient_db0` is the symbolic link to the raw partition `/dev/rfn_trans_db0` where the data is stored.

The backup software should back up the data in the character mode raw partition `/dev/rfn_trans_db0` either through links or the actual partition name. The software should also restore data from backups to the raw partition `/dev/rfn_trans_db0` either through links or the actual partition name. The software should also recreate the symbolic link during the restore.

FileNet Image Services for HP-UX

The `/fnsw/dev/1` directory contains character special files that point to the logical volumes in `/dev/<logical_volume_name>`.

The character special files in the `/fnsw/dev/1` directory have the same major and minor device numbers as the raw partition names the character special files point to in the `/dev/<logical_volume_name>` directory.

For example:

`/fnsw/dev/1/oracle_db0`

points to

`/dev/fnvg/rfn_oracle_db0`

where

`/fnsw/dev/1/oracle_db0` is the character special file containing the location of the raw partition `/dev/fnvg/rfn_oracle_db0` where the data is stored.

The backup software should back up the data in the character mode raw partition `/dev/fnvg/rfn_oracle_db0`, either through a character special file or the actual partition name. The software should also restore data from backups to the raw partition either through a character special file or the actual partition name. The software should also recreate the character special file during the restore.

FileNet Image Services for Solaris

The `/fnsw/dev/1` directory contains symbolic links to the raw slice or logical volume that contains the data. The first example shows the symbolic link to the logical volume name, when using logical volume management software (either VERITAS Volume Manager or Sun Enterprise Volume Manager). The second example shows the sym-

bolic link to the raw slice name for a FileNet system not managed by VERITAS.

Example for logical volumes with VERITAS Volume Manager:

```
/fnsw/dev/1/transient_db0
```

points to

```
/dev/vx/rdisk/fn_trans_db0
```

where:

/fnsw/dev/1/transient_db0 is the symbolic link to the raw partition (/dev/vx/rdisk/fn_trans_db0) in which the data is stored.

The backup software should back up the data in the character mode raw partition /dev/vx/rdisk/fn_trans_db0, either through links or the actual partition name. The software should also restore data from backups to the raw partition /dev/vx/rdisk/fn_trans_db0, either through links or the actual partition name. The software should also recreate the symbolic link during the restore.

Example for raw slices (without VERITAS Volume Manager):

```
/fnsw/dev/1/sec_db0
```

points to

```
/dev/rdisk/c0t0d0s6
```

where:

/fnsw/dev/1/sec_db0 is the symbolic link to the raw partition (/dev/rdisk/c0t0d0s6) in which the data is stored.

The backup software should back up the data in the character mode raw partition `/dev/rdisk/c0t0d0s6`, either through links or the actual partition name. The software should also restore data from backups to the raw partition `/dev/rdisk/c0t0d0s6`, either through links or the actual partition name. The software should also recreate the symbolic link during the restore.

Backing up and restoring `/fnsw/dev/1`

If your backup software can back up through links, you should be able to back up and restore through links in the `/fnsw/dev/1` directory. Otherwise, you must back up and restore using raw partition names. The restore program must be able to recreate the links from the backup, or the links must be recreated some other way.

Symbolic links

FileNet Image Services for AIX/6000 and FileNet Image Services for Solaris use symbolic links. The backup and restore procedure is:

- 1 Ensure that the symbolic links to the raw partitions are correctly set up.

See your platform-specific FileNet Image Services software installation procedures for additional information.
- 2 If your independent software vendor backup software can back up and restore raw partitions through symbolic links:
 - a Back up using the symbolic link names in the `/fnsw/dev/1` directory.
 - b Restore using the symbolic link names in the `/fnsw/dev/1` directory.

- 3 If your independent software vendor backup software can **not** back up and restore raw partitions through symbolic links:
 - a Back up using the actual raw partition names.
 - b Restore using the actual raw partition names.

Character special files

FileNet Image Services for HP-UX uses character special files. The backup and restore procedure is:

- 1 Ensure that the major and minor device numbers pointing to the raw partitions are correctly set up.

See your platform-specific FileNet Image Services software installation procedures for additional information.

- 2 Back up using either the character special file names in the `/fnsf/dev/1` directory or the actual raw partition names.
- 3 Restore using either the character special file names in the `/fnsf/dev/1` directory or the actual raw partition names.

Tip For a list of the symbolic link names and raw partition names, see **[“Raw partition and /fnsf/dev/1 names by server platform” on page 73.](#)**

No links

FileNet Image Services for Windows Server does not use links. The backup and restore procedure is simply:

- 1 Backup files in `\fnsf\dev\1`.

2 Restore files in \fnsw\dev\1.

Raw partition and /fnsw/dev/1 names by server platform

FileNet Image Services for AIX/6000

AIX logical volumes are always located in the /dev directory. This is the directory where you back up and restore the datasets containing the data. The datasets have the same names as the symbolic links in the /fnsw/dev/1 directory, but with an **rfn_** prefix, except the transient and permanent database names are shortened. The FileNet default dataset names are listed in /dev. However, these names could have been modified during installation.

| /fnsw/dev/1 directory | /dev directory |
|-----------------------|----------------|
|-----------------------|----------------|

MKF cache and database dataset names are:

| | |
|-----------------|-----------------|
| cache<n> | rfn_cache<n> |
| permanent_db<n> | rfn_perm_db<n> |
| permanent_rl<n> | rfn_perm_rl<n> |
| sec_db<n> | rfn_sec_db<n> |
| sec_rl<n> | rfn_sec_rl<n> |
| transient_db<n> | rfn_trans_db<n> |
| transient_rl<n> | rfn_trans_rl<n> |

The Oracle dataset names are:

| | |
|---------------|-------------------|
| oracle_db<n> | rfn_oracle_db<n> |
| oracle_rl<n> | rfn_oracle_rl<n> |
| oracle_sys<n> | rfn_oracle_sys<n> |
| oracle_udb<n> | rfn_oracle_udb<n> |
| oracle_tr<n> | rfn_oracle_tr<n> |

FileNet Image Services for HP-UX

On HP-UX servers, raw partition names are located under /dev in a subdirectory with the volume group name /dev/<volume_group>. For example, if the volume group name is fnvg, the path to the logical volumes is /dev/fnvg.

The logical volume datasets have the same names as the dataset names in the /fnsw/dev/1 directory, but with an **rfn_** prefix, except that transient and permanent databases names are shortened. The datasets in /fnsw/dev/1 and the datasets in /dev/fnvg are considered character special files. Each character special file name under /fnsw/dev/1 (such as cache0) has the same major device and minor device numbers as the corresponding character special file name under /dev/<volume_group> (such as rfn_cache0). The same major device and minor device numbers point to the same raw partition.

The FileNet default raw partition names are listed in /dev/<volume_group>. However they could have been modified during installation.

| /fnsw/dev/1 directory | /dev/<volume_group> |
|---|---------------------|
| MKF cache and database dataset names are: | |
| cache<n> | rfn_cache<n> |
| permanent_db<n> | rfn_perm_db<n> |
| permanent_rl<n> | rfn_perm_rl<n> |
| sec_db<n> | rfn_sec_db<n> |
| sec_rl<n> | rfn_sec_rl<n> |
| transient_db<n> | rfn_trans_db<n> |
| transient_rl<n> | rfn_trans_rl<n> |

The Oracle dataset names are:

| | |
|---------------|-------------------|
| oracle_db<n> | rfn_oracle_db<n> |
| oracle_rl<n> | rfn_oracle_rl<n> |
| oracle_sys0 | rfn_oracle_sys0 |
| oracle_udb<n> | rfn_oracle_udb<n> |

`oracle_tr<n>``rfn_oracle_tr<n>`

FileNet Image Services for Solaris

Two techniques for dataset configuration are available on the Solaris platform.

- a Slices are managed by Solaris native tools (IDMIS 3.3.1 and earlier). Slices are similar to UNIX partitions and have long path names.
- b Logical volumes are managed by volume management software: either VERITAS, or Sun Enterprise Volume Manager.

Logical volumes are located under `/dev/vx/rdisk/<vol_grp>` where `<vol_grp>` is the volume group name.

Unlike FileNet Image Services for AIX/6000 and FileNet Image Services for HP-UX, the `rdsk` directory indicates that the logical volumes are raw logical volumes. Therefore no `r` is prefixed to the logical volume name. The names in `/fnsd/dev/1` are symbolic links to the logical volume in the `/dev/vx/rdisk` directory.

Tip Since FileNet Image Services for HP-UX and FileNet Image Services for Solaris do not have the name length restriction for logical volume names as FileNet Image Services for AIX/6000 has, the words permanent and transient may be spelled out, for example, `rfn_permanent_db0` instead of `rfn_perm_db0`.

FileNet default dataset names are listed under `/dev/vx/rdisk/<volume_group>`. However these names could have been modified during installation.

| /fnsw/dev/1 directory | /dev/vx/rdsk/<volume_group> directory |
|---|--|
| MKF cache and database dataset names are: | |
| cache<n> | fn_cache<n> |
| permanent_db<n> | fn_perm_db<n> |
| permanent_rl<n> | fn_perm_rl<n> |
| sec_db<n> | fn_sec_db<n> |
| sec_rl<n> | fn_sec_rl<n> |
| transient_db<n> | fn_trans_db<n> |
| transient_rl<n> | fn_trans_rl<n> |

The Oracle dataset names are:

| | |
|---------------|------------------|
| oracle_db<n> | fn_oracle_db<n> |
| oracle_rl<n> | fn_oracle_rl<n> |
| oracle_udb<n> | fn_oracle_udb<n> |
| oracle_sys0 | fn_oracle_sys0 |
| oracle_tr<n> | fn_oracle_tr<n> |

Finding raw partition names

You need to know the location of the raw partitions to check that the links are correct or to rebuild the links. You can use the following methods to find the link directory name or actual raw partition name where the data is located. Your raw partition names may be different from the ones in the examples, depending on how your system is configured.

FileNet Image Services for AIX/6000

- 1 Change the current directory to /fnsw/dev/1.
- 2 Enter the **ls -l** command to list all symbolic links.

The following is a sample screen output:

```
0 lrwxrwxrwx 1 fnsf fnusr 15 May 31 10:28 cache0@ -> /dev/rfn_cache0
0 lrwxrwxrwx 1 fnsf fnusr 18 May 31 10:28 odda1@ -> /dev/fnsod.0,0,1,1
0 lrwxrwxrwx 1 fnsf fnusr 0 May 31 10:28 osara@ -> /
0 lrwxrwxrwx 1 fnsf fnusr 17 May 31 10:28 permanent_db0@ ->/dev/rfn_perm_db0
0 lrwxrwxrwx 1 fnsf fnusr 17 May 31 10:28 permanent_db1@ ->/dev/rfn_perm_db1
0 lrwxrwxrwx 1 fnsf fnusr 17 May 31 10:28 permanent_db2@ ->/dev/rfn_perm_db2
0 lrwxrwxrwx 1 fnsf fnusr 17 May 31 10:28 permanent_rl0@ ->/dev/rfn_perm_rl0
0 lrwxrwxrwx 1 fnsf fnusr 16 May 31 10:28 sec_db0@ -> /dev/rfn_sec_db0
0 lrwxrwxrwx 1 fnsf fnusr 16 May 31 10:28 sec_db1@ -> /dev/rfn_sec_db1
0 lrwxrwxrwx 1 fnsf fnusr 16 May 31 10:28 sec_db2@ -> /dev/rfn_sec_db2
0 lrwxrwxrwx 1 fnsf fnusr 16 May 31 10:28 sec_rl0@ -> /dev/rfn_sec_rl0
0 lrwxrwxrwx 1 fnsf fnusr 19 May 31 10:28 oracle_db0@ ->/dev/rfn_oracle_db0
0 lrwxrwxrwx 1 fnsf fnusr 20 May 31 10:28 oracle_rl0@ ->/dev/rfn_oracle_rl0
```

FileNet Image Services for HP-UX

- 1 Change the current directory to /fnsf/dev/1.
- 2 Enter the **ls -l** command to find the major and minor device file numbers for each character special file. In the example below, the major device file number is 64 and the minor device file number is 0x010005 for the Oracle index database.

Example Output:

```
crw-rw-r-- 1 oracle dba 64 0x010005 Jun 4 11:40 oracle_db0
```

- 3 For each character special file name under /fnsf/dev/1 (such as oracle_db0), make sure that the corresponding character special file name under /dev/fnvg (such as rfn_oracle_db0) has the same major device number and minor device number. The same major device and

minor device numbers point to the same raw partition and define the location of the raw partition.

In each character special file directory, enter the **ls -l** command to find the major and minor device numbers. For example, enter the following command to display the Oracle database link:

ls -l rfn_oracle_db0

Example Output:

```
crw-rw-r-- 1 oracle dba 64 0x010005 Jun 4 11:47 /dev/fnvg/rfn_oracle_db0
```

The example above shows a major number of 64 and a minor number of 0x010005 for the index database logical volume. Note these numbers and compare them to the major and minor numbers in step 2. The major device and minor device numbers for rfn_oracle_db0 must be the same as the major device and minor device numbers for oracle_db0.

FileNet Image Services for Solaris

Use one of two methods, depending on whether your system is configured in raw slices or logical volumes.

When your system is configured in slices:

- 1 Change the current directory to /fnsd/dev/1 directory.
- 2 Enter the **ls -l** command.

Example Output:

```
lrwxrwxrwx 1 root other 18 May 10 2001 oracle_db0 --> /dev/rdsd/c0t1d0s1
```

3 Change to the /dev/rdsk directory.

4 Enter the following command:

ls -l c0t1d0s1

Example Output:

```
lrwxrwxrwx 1 root root 88 May 10 2001 c0t1d0s1 -->../../devices
/iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0:b,raw
```

The path to the oracle_db0 raw slice in this example is:

```
/devices/iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/
sd@1,0:b,raw
```

When your logical volumes are managed by a logical volume management product (either VERITAS or Sun Enterprise Volume Manager:

1 Change the current directory to /fnsw/dev/1.

2 Enter the **ls -l** command.

Example Output:

```
lrwxrwxrwx 1 root other 31 May 16 13:20 oracle_db0 -->
/dev/vx/rdsk/fnvg/fn_oracle_db0
```

The path to the logical volume in this example is:

`/dev/vx/rdsk/fnvg/fn_oracle_db0`

Advanced backup methods

Use Advanced Backup methods if you wish to back up and restore individual raw datasets and your independent software vendor software has the capability to do so.

Attention

If you use Advanced Backup, you must use an Advanced Restore. If you use Advanced Backup, you cannot use a Simple Restore.

- 1 Stop the FileNet Image Services software by doing one of the following. You must be a member of the fnadmin or fnop and fnusr groups to stop/start FileNet Image Services software.
 - On FileNet Image Services for UNIX platforms, enter:
Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)
 - On FileNet Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
 - Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line.

For FileNet-controlled databases, always use the Task Manager or command line to stop both FileNet Image Services software and the RDBMS database. Do **not** use the Microsoft SQL Service Manager to stop the RDBMS. If you do, an error displays because the Microsoft SQL Service is dependent on the IS ControlService.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

The Stop button in the Task Manager or the `initfnsw stop` command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **`killfnsw`** command, described in Appendix B under **[“killfnsw” on page 171](#)**, then run the `initfnsw stop` command.

Important

Rebooting your operating system might automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the `initfnsw stop` command or the Stop button in Task Manager before you back up or restore.

- 2 Perform the backup using your independent software vendor backup software. Back up the FileNet file systems, databases and cache depending on the backup schedule.

If the FileNet Image Services software on your platform uses symbolic links or character special files, you may need to employ special procedures to back up the data referenced in the `/fnsw/dev/1` directory. This depends on whether your independent software vendor backup software is capable of backing up raw partitions through links. Refer to **[“Backing up and restoring /fnsw/dev/1” on page 71](#)**.

For a list of FileNet and non-FileNet Image Services software to back up, refer to **[“File System Software and Data to Back Up” on page 50](#)**.

- 3 Start up the FileNet Image Services software using the method appropriate for your platform:

Important

Start the RDBMS database first, if you have a site-controlled RDBMS database, **before** starting FileNet Image Services software.

- On FileNet Image Services for AIX/6000, HP-UX, and Solaris, enter:

Xtaskman &

Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)

- On FileNet Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start the FileNet Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfns -y start** at the command line.

Advanced restore methods

This chapter describes Advanced Restore methods that an experienced and trained system administrator can use to restore a FileNet system with independent software vendor restore software packages. The chapter provides you with:

- A summary overview table describing restore and recovery tasks by type of datasets lost.
- A table describing types of data recovery methods by database to be restored.
- Guidelines for rebuilding your total system, if necessary.
- Procedures to restore by type of dataset lost.
- Multiserver restore information.

Important

You are responsible for your restore operation if you use independent software vendor backup software.

To use Advanced Restore methods, your independent software vendor software must be able to restore raw partitions. You can use Advanced Restore methods in cases where you only need to restore a single database or individual datasets or want to use more specialized data recovery methods.

Advanced Restore methods are intended to:

- Restore your MKF databases, index database, and cache to the same point in time as the end of your last backup.
- Recover or re-enter any updates up to the moment of the system failure.

If you have mirrored all your disks, a restore is necessary only if both spindles of a mirrored pair of disks are lost. This is a rare occurrence.

Attention

To ensure that you can restore your backups, it is imperative that you perform a **practice backup and restore** on a test system. Verify that your restore software works, especially for MKF and index database restores. See **“Testing third-party backup/restore” on page 22** for more information.

In addition, you should backup your Image Services datasets every six months using FileNet Enterprise Backup/Restore (ERB) software. This will ensure that the foundation of your databases can be rebuilt from a method recognized by FileNet. Maintaining reasonably current EBR backups could reduce the length of recovery time of your databases. Store your FileNet backups either at an off-site location or in a fireproof safe

Summary of advanced restore task order

The following summarizes the Advanced Restore tasks by type of dataset lost and by sequential order of the task performed. Please reference **“Table 4-1: Restore and Recovery Tasks for Single and Multiple Lost Datasets” on page 86** throughout this discussion. In this table, a Y(es) or N(o) indicates whether the task applies to the dataset to be restored.

Single dataset restore task order

The restore and data recovery/re-entry tasks for a single dataset are summarized in the applicable row in [“Table 4-1: Restore and Recovery Tasks for Single and Multiple Lost Datasets” on page 86](#). These main tasks are described in detailed procedures in [“Restore methods by dataset type” on page 89](#).

Multiple dataset restore task order

If you lose more than one dataset, you need to determine a task order in which you restore these datasets. One method to order these tasks is to go **column by column** in the chart [“Table 4-1: Restore and Recovery Tasks for Single and Multiple Lost Datasets” on page 86](#). You perform all applicable tasks in one column before moving to the next column of tasks. An alternative method to order the recovery tasks is to simply perform the tasks **row by row** for each lost dataset.

Multiple dataset restore examples

If you lost the permanent database and transient database, you can do the following: perform the initialize task for the transient database in column 1 first before restoring the permanent database as indicated in column 2.

If you lost the permanent database and the index database, sequentially perform all applicable tasks in the permanent database row and then continue to perform all tasks in the index database row. The tasks are described in detailed procedures in [“Restore methods by dataset type” on page 89](#).

Table 4-1: Restore and Recovery Tasks for Single and Multiple Lost Datasets

| Lost Dataset | Restore and Data Recovery Tasks (in Sequential Order) | | | | | | |
|---|---|----------------|--|------------------------|-------------|-------------|-----------------------------------|
| | 1 Initialize | 2 Restore | 3 Roll- forward (MKF or RDBMS) | 4 Reset Scalar # | 5 Import | 6 Rescan | 7 Manually Re-enter Data |
| /fnswh/local (UNIX) or \\fnswh_loc (Windows Server) | N | Y ^a | N | N | N | N | Y ^a |
| Permanent database | N | Y | Y | Y | Y | N | N |
| Permanent database re- covery log | Y ^b | N | N | N | N | N | N |
| Security database | N | Y | Y | N | N | N | Y |
| Security database recov- ery log | Y ^b | N | N | N | N | N | N |
| Index database | N | Y | Y | N | Y | N | Y ^c |
| RDBMS redo/transaction log | N | Y | N | N | N | N | N |
| WorkFlo queue database | N | Y | Y | N | Y | N | N |
| WorkFlo queue recovery log | N | Y ^d | Y | N | Y | N | N |
| Transient database and cache (if restoring) | Y ^e | Y | N | N | N | Y | N |
| Transient database and cache (if not restoring) | Y ^e | N | N | N | N | Y | N |
| Transient database recov- ery log | Y ^b | N | N | N | N | N | N |

- a. The /fnswh/local (or \fnswh_loc on Windows Server) file system contains the FileNet configuration files. After a restore of this file system, if you have changed your FileNet configuration database (CDB) file since the last backup, you must copy the restored CDB file into the next highest numbered CDB file. The restored system state should match the current CDB file used by the Configuration Editor.

The CDB file number is advanced to the next highest number each time you update your FileNet system configuration. (Older files are not automatically deleted.) When you restart your system after a full restore, the Configuration Editor uses the highest numbered CDB file found in the directory.

For example, at the time you performed the last full backup, the CDB file number was 10 (IMS_10.cdb). Later you updated the configuration to add extents and datasets and the current CDB file is 11 (IMS_11.cdb). After a full restore, the CDB file that correctly corresponds to the backup date is IMS_10.cdb, not IMS_11.cdb. Copy IMS_10.cdb into the next highest numbered file IMS_12.cdb.

The CDB file is found in:

/fnswh/local/sd/conf_db/IMS_nnn.cdb (UNIX)

<drive>:\fnswh_loc\sd\conf_db\IMS_nnn.cdb (Windows Server)

After copying the restored CDB file to the next highest numbered file, use the fn_edit utility or Configuration Editor and re-enter changes made to the configuration since the last backup. For example, the last backup was done when IMS_10.cdb was created. A new dataset was added in IMS_11.cdb. When the user copies IMS_10.cdb to IMS_12.cdb, as in the example, the new dataset configuration change is lost and must be re-entered.

- b. Initialize the recovery log using the MKF_zeroaj system tool. Do not use the fn_util initperm, fn_util inittrans, or fn_util initsec system tool. If MKF_zeroaj fails, perform the restore and data recovery tasks in sequential order from the above table for the affected database. This failure means the recovery log has lost data needed by the database to maintain data integrity. Therefore the database data is in an unknown state and the database should be restored.
- c. Reindex documents if you usually index documents after committal or change indexes after committal. You do not need to reindex any documents if your indexes are complete and correct at time of committal.
- d. To restore only the recovery log, refer to your RDBMS documents for procedures. If you cannot restore the recovery log, then restore both the WorkFlo Queue database and the recovery log.
- e. Do not initialize if you have a cache-only system unless you have reliable backups of the transient database and cache.

Requirements for a restore

Depending on what was lost (logical volume groups, logical volumes, file systems, and/or links), you may have to rebuild your system before attempting to restore any FileNet datasets. To rebuild your system prior to a restore, perform the following steps where necessary, **in the order given**:

- 1 Restore the operating system if it is lost.
- 2 Create your logical volume groups.
- 3 Create your logical volumes.

Image Services for UNIX requires character mode special files, not block mode special files. If your operating system tool recreates both types, be sure to link to the character mode special file.

- 4 Create file systems for FileNet and non-FileNet software.
- 5 Restore any corrupted file systems. For example, restore the following FileNet file systems, if needed, before restoring the datasets:

- /fnsf (or \fnsf on Image Services for Windows Server)
- /fnsf/local (or \fnsf_loc on Image Services for Windows Server)

These file systems contain the FileNet Image Services software configuration required by the FileNet databases. The /fnsf/dev/1 directory contains symbolic links or character special files pointing to data located in raw partitions. Links or character special files must be restored before you restore the data.

FileNet Image Services software uses symbolic links or character special files depending on platform. Whether your independent software vendor backup software is capable of restoring raw partitions through

links or special files raises issues on how to restore data referenced in the

/fnsw/dev/1 directory. For a description of how to do this, see [**“Backing up and restoring /fnsw/dev/1” on page 71.**](#)

Important

On an Image Services for AIX/6000 system, you can use your mksysb tape to restore the operating system. Your FileNet file systems may also be on the mksysb tape if they were part of the rootvg volume group.

- 6 Create the symbolic links or character special files in /fnsw/dev/1, if any are used.
- 7 Restore lost datasets using your independent software vendor software. Refer to the appropriate procedures in this chapter.

For more information on how to create logical volume groups, logical volumes, file systems, links and character device files, see your platform-specific Image Services software installation procedures.

Restore methods by dataset type

This section describes the advanced restore methods by the type of dataset lost. If you have lost more than one dataset, you can refer to the table, [**“Table 4-1: Restore and Recovery Tasks for Single and Multiple Lost Datasets” on page 86,**](#) for a roadmap on how to proceed. Alternatively, you can perform each set of procedures in this section for each lost dataset.

For expanded explanations pertaining to the scalar numbers table, the import process, MKF rollforward recovery, and synchronization issues, see [**Chapter 5, “Synchronization and recovery information,” on page 139.**](#)

Attention

If you use Advanced Backup methods, you must use an Advanced Restore methods.

Permanent database

Use this method to restore a lost permanent database.

- 1 Stop the FileNet Image Services software by doing the following. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. Do **not** use the Microsoft SQL Service Manager to stop the RDBMS. If you do, an error displays because the Microsoft SQL Service is dependent on the Image Services ControlService.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

The Task Manager or the `initfnsw stop` command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in Appendix B under **“killfnsw” on page 171**, then run the `initfnsw stop` command.

Attention

Rebooting your operating system may automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the `initfnsw stop` command or the Stop button in Task Manager before you back up or restore.

- 2 Make a copy of the current disk-resident scalar numbers table checkpoint file, `/fnsw/local/sd/snt.chkpt`, before restoring `/fnsw` or `/fnsw/local` (`/fnsw` and `\fnsw_loc` on Image Services for Windows Server).

Copy the file to tape or to a directory such as `/fnsw/local/tmp`. You must save a copy of the current checkpoint file at this point because you need the file to update the scalar numbers table, described in **Step 6 on page 92**.

- 3 Before restoring any data, make sure you have backed up the current MKF recovery logs for the permanent database provided that the current recovery logs are not corrupt.

This precaution ensures that you have the current recovery logs if you unintentionally restored the old recovery logs.

Attention

Do not restore any MKF recovery logs when you restore the MKF database (The exception is when you have lost both the MKF database and the recovery log). The current recovery logs contain all

changes up to the moment of the crash - you will want to roll those forward.

Remember to back up your current permanent database recovery logs **before** doing a restore. This precaution provides protection should your first restore fail.

- 4 Perform a restore of the permanent database from tapes made with your independent software vendor backup software.
- 5 Start the FileNet databases (permanent, transient, security, and index databases) using the `fn_util startdb` Image Services command. You must be a member of the `fnadmin` or `fnop` and `fnusr` and `dba` groups to do this.

Run the command from the system console on each Image Services server. Enter the following at the server command line:

fn_util startdb

Although the databases are started, FileNet Image Services software is still stopped. This is to prevent users from entering documents at this recovery phase.

- 6 Update the scalar numbers table with a valid checkpoint file.


Important

Advancing the document ID number in the scalar numbers table is critical to maintaining data integrity. This will prevent the system from creating duplicate document ID numbers in the permanent database as new documents are entered. Duplicate document ID numbers can lead to unpredictable results.

To advance the scalar numbers table, do the following:

- a Verify you have a current checkpoint file (snt.chkpt) with which to update the scalar numbers table. As long as /fnsw/local has not been restored and is not corrupted, the checkpoint file is valid and current. The checkpoint file, located in /fnsw/local/sd/snt.chkpt, is always **invalid** if you have had to restore /fnsw/local as part of your overall restoration operation.
- b Copy the **saved** version (not a **restored** version) of the checkpoint file back to /fnsw/local/sd if you have a recently saved copy.
- c If you could not save the checkpoint file or it has been corrupted by a disk crash, perform both of the following tasks:
 - Estimate how much to advance the document ID number in the snt.chkpt file.

For example, if you typically scan 20,000 documents per day and your backup is two days old, you should estimate that your document ID number would increase by two days' worth of documents. Then add one extra day's worth of document IDs to this number (for example, 40,000 + 20,000) to ensure that you do not underestimate. This means you should advance the document ID number by 60,000. It is better to increase by too much than by too little.
 - Edit the scalar numbers table snt.chkpt file to advance the document ID number.
 - Change to the /fnsw/local/sd directory (or \fnsw_loc\sd directory on Image Services for Windows Server).
 - Modify the document ID number using an appropriate platform-specific editor such as vi for UNIX or edit or WordPad for Windows Server. The editor displays a screen resembling the following:



- Save the file after modifying the document ID number.
- The fourth number represents the date and time of the last update to the snt.chkpt file and translates into a time stamp when you run the SNT_update program.

d Run the SNT_update program, using the following procedures:

- Log on as a FileNet user who is a member of the fnadmin or fnop and fnusr groups.
- Enter the command: **SNT_update**
- Enter **Y** at the prompt.

SNT_update displays messages indicating either successful completion or errors that prevented successful update of the database. For a detailed description of the SNT_update program, see [**“Updating the scalar numbers table” on page 146.**](#)

- 7 Restart FileNet Image Services software using the appropriate method for your platform. Restart stops and starts FileNet Image Services software.

Important

If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Restart** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Restart** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y restart** at the command line.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start up the Image Services software.

Important

When FileNet Image Services software is started, MKF database software automatically rolls forward the permanent database.

- 8 Update the optical media surface ID number in the checkpoint.osa file. Use the Identify Media in Library option of the Storage Library Control Program to get this number.

The optical media surface ID number is in the checkpoint.osa data file, located in the /fnsw/local/sd (or \fnsw_loc\sd) directory. The program locates all known optical media and automatically updates the optical media surface ID number in the checkpoint.osa file with the new numbers.

This step provides the input for the next step, which updates the surface ID number in the scalar numbers table snt.chkpt file.

- 9 Manually update the optical media surface ID number in the snt.chkpt file using input from the DOC_tool program.

The surface ID number must be updated **before** document entry resumes to prevent duplicate surface ID numbers on the next new optical media.

- a At a command line or MS-DOS prompt, enter:

DOC_tool

- b At the DOC_tool prompt, enter:

dev

The dev subcommand displays the contents and surface ID for the gripper and each drive and slot in the storage library from the checkpoint.osa file.

Respond as follows to the prompts:

Print status of empty slots? (y/n) : **n**

Print status of disabled slots? (y/n) : **y**

Your display will look similar to:

| | | | |
|---------|----|----------------|-------------|
| Gripper | | Contents=known | Surfid=4002 |
| Drive | 0 | Contents=known | Surfid=4000 |
| Drive | 1 | Contents=known | Surfid=3998 |
| Slot | 21 | Contents=known | Surfid=3996 |
| Slot | 28 | Contents=known | Surfid=3994 |

- c Find the highest numbered surface ID from your display. Compare this number to the surface ID number in the snt.chkpt file.

- d If the surface ID in the snt.chkpt file is not two numbers higher than the surface ID displayed in DOC_tool, edit the snt.chkpt file to increase the surface ID number by two.

For example, if the surface ID number in the snt.chkpt file is 4000 and DOC_tool displays a surface ID number of 4002, then manually update the surface ID number in snt.chkpt to 4004.

- e Rerun the SNT_update program by entering:

SNT_update

- f Enter **Y** when prompted to update the scalar numbers table.

- 10 Import documents from optical media using the procedures which follow. This will further ensure that any migrated documents lost between the backup and the point of failure will be recovered to the permanent database.

Attention

If you do not have a current checkpoint file (/fnsw/local/sd/snt.chkpt), from which to update the scalar numbers table, you must use either the odrecover or stdocimp (with the updatesnt option) utility to advance the scalar numbers table **before you enter any documents**.

- a Import both the primary media and the tranlog media that have been written to since the backup was performed. Use one of the following FileNet import utilities:
 - If you have defined WorkFlo queues for each document class and you want documents reinserted into WorkFlo queues, use **Import documents from media** from the Background Job Control menu.

- If you want more flexible options in importing documents, use the **odrecover** or **stdocimp** utility. The **odrecover** utility does not import documents into WorkFlo queues. Although the **stdocimp** utility imports by default to WorkFlo queues, it allows you the option to not insert documents into a WorkFlo queue if a queue is defined for the document class.
- b You do not need to identify the individual lost documents or the range of documents IDs to import, as long as you can determine which optical media have been written to most recently. The import programs add missing documents to either the index database or permanent database or both. The import programs do not create duplicate database entries.

Although tranlog media do not usually need to be imported, it is safer to import recently written tranlog media, in the event the disk failure occurred before the tranlog media information was completely written to the primary media. Recently written tranlog media are those written to after the backup was taken.

Tip If you maintain a daily log of optical media surfaces which have been written to, you can easily determine which optical media to import.

Refer to the *Image Services System Administrator's Handbook* for more information on the import methods or the **stdocimp** utility. To download IBM FileNet documentation from the IBM support page, see [**“Accessing IBM FileNet documentation” on page 14.**](#)

Refer to [**“odrecover” on page 176**](#) for a description of the **odrecover** utility.

- 11 At the next convenient non-production time, run the **dbverify** tool to determine if erroneous document IDs exist in the permanent database.

The dbverify tool compares document IDs in the permanent database against the index database.

When you run dbverify, you should select only the range of document IDs affected. The program runs faster when you limit the range and do not select a comparison of all document IDs in the permanent database and index database.

To access online help information for using dbverify, enter from a command line prompt or MS-DOS command prompt:

dbverify -h

Permanent database recovery log

Use these steps to clear data in a lost or corrupt permanent database recovery log. This allows the recovery log to be used successfully again by the permanent database.

When MKF database software starts, it determines if a rollforward of the recovery log is necessary. Rolling forward a corrupt recovery log will result in a fatal error condition and prevent Image Services from starting. A cleared recovery log means there is no data to roll forward to the permanent database.

These steps do **not** perform a recovery log restore and assume:

- The permanent database is intact and has not been restored.
- Only the permanent database recovery log is lost or corrupted.
- No current backup of the recovery log exists.

Procedures are:

- 1 Stop the FileNet Image Services software using one of the following methods. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the Microsoft SQL Service is dependent on the Image Services ControlService.

The Task Manager or the initfnsw stop command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure initfnsw stop is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [**“killfnsw” on page 171**](#), then run the initfnsw stop command.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

2 To clear data in or initialize the recovery log, do the following:

- a Run the **MKF_zeroaij** utility on the permanent database from a command line prompt or MS-DOS command prompt.

MKF_zeroaij zeroes out the specified MKF database recovery log. Running MKF_zeroaij prevents a corrupted or lost recovery log from being applied (or rolled forward) to the permanent database when the permanent database next opens.

See **“MKF_zeroaij” on page 174** for a description of the MKF_zeroaij syntax.

- b If you receive a fatal error message, go to step c.

If you receive no error message, skip to step 3.

- c Restore the permanent database if you receive a fatal error message.

If you receive the following fatal error message while attempting to run MKF_zeroaij, MKF_zeroaij has detected that the recovery log has lost data needed by the permanent database to maintain data integrity:

FATAL ERROR: information in the recovery log exists that is required for database integrity. The system cannot allow the recovery log (formerly called the aij, or ‘after image journal’) to be zeroed.

The resolution to the fatal error is to restore the permanent database. Refer to **“Permanent database” on page 90** for restore procedures.

- 3 Start the software using the method appropriate for your platform:

Important

Start the RDBMS database first, if you have a site-controlled RDBMS database, **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y start** at the command line or the MS-DOS command prompt.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start the Image Services software.

- 4 If only the permanent database recovery log was corrupted or lost (not the permanent database), the permanent database is now ready for production.

Security database

Use these steps to restore a lost security database.

- 1 Stop the FileNet Image Services software by doing the following. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.

- On Image Services for UNIX servers, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager or the command line to stop both FileNet Image Services software and the RDBMS database. Do **not** use the Microsoft SQL Service Manager to stop the RDBMS. If you do, an error displays because the Microsoft SQL Service is dependent on the Image Services ControlService.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

Using the Task Manager or the initfnsw stop command is the normal method to stop the system. If normal attempts fail to stop the system,

evaluate your software state. Be sure `initfnsn stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **`killfnsn`** command, described in [“**killfnsn**” on page 171](#), then run the `initfnsn stop` command.

Important

Restarting your operating system might automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the `initfnsn stop` command or the Stop button in Task Manager before you back up or restore.

- 2 Before restoring any data, make sure you have backed up the current MKF recovery logs for the security database provided that the recovery logs are current and not corrupted.

This precaution ensures that you have the current recovery logs if you unintentionally restored the old recovery logs or if initial recovery efforts should fail.

- 3 Perform a restore of the security database from tapes made with your independent software vendor backup software.

Important

Do not restore any MKF recovery logs when you restore the MKF database. (The exception is when you have lost both the permanent database and the permanent database recovery log.) The current recovery logs contain all changes up to the moment of the crash - you will want to roll those forward.

Remember to back up your current security database recovery logs **before** doing a restore. This precaution provides protection should your first restore fail.

- 4 Restart FileNet Image Services software using the method appropriate for your platform. You must be a member of the fnadmin or fnop and fnusr and dba groups to do this.

Important

If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Restart** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Restart** button. Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start up the Image Services software.
 - Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y restart** at the command line or the MS-DOS command prompt.
- 5 Manually re-enter any missing data that was not on the backup tape. For example, you may have to re-enter users, user groups, and other security information added or deleted since the backup was taken.

Security database recovery log

Use these steps to clear data in a lost or corrupt security database recovery log. This will allow the recovery log to be used successfully again by the security database.

When MKF database software starts, it determines if a rollforward of the recovery log is necessary. Rolling forward a corrupt recovery log will result in a fatal error condition and prevent Image Services from starting. A cleared recovery log means there is no data to roll forward to the permanent database.

This method does **not** perform a recovery log restore and assumes:

- The security database is intact and has not been restored.
- Only the security database recovery log is lost or corrupted.
- No current backup of the recovery log exists.

Procedures are:

- 1 Stop the FileNet Image Services software using the following method. You must be a member of the fnadmin or fnop and fnusr groups to stop/start Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS data-

base. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

Using the Task Manager or the `initfnsw stop` command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [“killfnsw” on page 171](#), then run the `initfnsw stop` command.

- 2 To clear data in or initialize the recovery log, do the following:
 - a Run the **MKF_zeroaij** utility on the security database from a command line prompt or MS-DOS command prompt.

MKF_zeroaij zeroes out the MKF database recovery logs. Because your recovery log is corrupted or lost, running **MKF_zeroaij** prevents bad recovery log data from being applied (or rolled forward) to the security database when the security database next opens.

See [“MKF_zeroaij” on page 174](#) for a description of the **MKF_zeroaij** syntax.

- b If you receive a fatal error message, go to step c.

If you receive no error message, skip to step 3.

- c Restore the security database if you receive a fatal error message.

If you receive the following fatal error message while attempting to run MKF_zeroaij, MKF_zeroaij has detected that the recovery log has lost data needed by the permanent database to maintain data integrity:

FATAL ERROR: information in the recovery log exists that is required for database integrity. The system cannot allow the recovery log (formerly called the aij, or 'after image journal') to be zeroed.

To resolve the fatal error, restore the security database. Refer to [**"Security database" on page 103**](#) for restore procedures.

- 3 Start the software using the appropriate method for your platform:

Important

Start the RDBMS database first, if you have a site-controlled RDBMS database, **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y start** at the command line or the MS-DOS command prompt.
- 4 If only the security database recovery log was corrupted or lost (not the security database), the security database is now ready for production.

Index database

Use these steps to restore a lost index database.

- 1 Stop the FileNet Image Services software by using one of the following methods. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line or the MS-DOS command prompt.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

Using the Task Manager or the **initfnsw stop** command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure **initfnsw stop** is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [“killfnsw” on page 171](#), then run the **initfnsw stop** command.

Important

Rebooting your operating system may automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the **initfnsw stop** command or the Stop button in Task Manager before you back up or restore.

- 2 If your RDBMS (Oracle or Microsoft SQL Server) software needs to be restored, restore that before you restore the index database.

See your RDBMS documentation or independent software vendor product documentation for restore procedures.

To restore RDBMS software, such as Oracle control files, you need read and write permissions on RDBMS files. You need to be a member of the dba group.

- 3 Perform a restore of the index database from tapes made with your independent software vendor backup software.

- 4 Restart FileNet Image Services software using the method appropriate for your platform. Restart shuts down and starts FileNet Image Services software.

Important

Start the RDBMS first, if you have a site-controlled RDBMS database, **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Restart** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Restart** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start up the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y restart** at the command line or the MS-DOS command prompt.

- 5 Apply the index database redo or transaction log(s) to roll the index database forward.

Refer to your RDBMS documentation for instructions on recovering and applying redo or transaction logs.

- 6 Import documents from optical media using the following procedures. This will further ensure that any migrated documents lost between the backup and the point of failure will be recovered to the index database.

- a Import both the primary media and the tranlog media that have been written to since the backup was performed. Use one of the following FileNet import utilities:
- If you have defined WorkFlo queues for each document class and you want documents reinserted into WorkFlo queues, use the “**Import documents from media**” from the Background Job Control menu.
 - If you want more flexible options in importing documents, use the **odrecover** or **stdocimp** utility. The **odrecover** utility does not import documents into WorkFlo queues. Although the **stdocimp** utility imports by default to WorkFlo queues, it allows you the option to not insert documents into a WorkFlo queue if a queue is defined for the document class.

You do not need to identify the individual lost documents or the range of documents IDs to import, as long as you can determine which optical media should be recovered. These would be the optical media that have been written to most recently.

The import programs add missing documents to either the index database or permanent database or both. The import programs do not create duplicate database entries.

Although tranlog media do not usually need to be imported, it is safer to import recently written tranlog media, in the event the disk failure occurred before the tranlog media information was completely written to the primary media. Recently written tranlog media are those written to after the backup was taken.

Tip

If you maintain a daily log of optical media surfaces which have been written to, you can easily determine which optical media to import.

Refer to the *Image Services System Administrator's Handbook* for more information on the import methods or the stdocimp utility. To download IBM FileNet documentation from the IBM support page, see [**“Accessing IBM FileNet documentation” on page 14.**](#)

Refer to [**“odrecover” on page 176**](#) for a description of the odrecover utility.

- 7 Reindex the imported documents by manually retyping the entries, if you typically change indexes after committal. The imported documents are imported with their originally indexed entries.
- 8 At the next convenient non-production time, run the **dbverify** tool to determine if erroneous document IDs exist in the index database. The dbverify tool compares document IDs in the index database against the permanent database.

when you run dbverify, you should select only the range of document IDs affected. The program runs faster when you limit the range and do not ask for a comparison of all document IDs in the permanent and index databases.

To access online help information for using dbverify, enter from a command line prompt or MS-DOS command prompt:

dbverify -h

RDBMS redo/transaction log

To restore a lost or corrupt Oracle redo log or Microsoft SQL Server transaction log, refer to your RDBMS system administration documentation. Restore procedures for corrupted RDBMS logs are beyond the scope of this manual.

WorkFlo Queue database

Use this method to restore a lost WorkFlo Queue database. A WorkFlo Queue database contains WorkFlo queues tables and is managed by Oracle or Microsoft SQL Server. You have a WorkFlo Queue database if you have a dedicated application server running WorkFlo Queue services.

- 1 Stop the FileNet Image Services software by doing one of the following. You must be a member of the fnadmin or fnop and fnusr groups to stop/start Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y start** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager or command line to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

Important If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

Using the Task Manager or the `initfnsw` stop command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw` stop is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **`killfnsw`** command, described in [“`killfnsw`” on page 171](#), then run the `initfnsw` stop command.

Important Rebooting your operating system may automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the Stop button in Task Manager before you back up or restore.

- 2 If your RDBMS (Oracle or Microsoft SQL Server) software needs to be restored, restore that before you restore the index database.

See your RDBMS documentation or independent software vendor product documentation for restore procedures.

- 3 Perform a restore of the WorkFlo Queue database from tapes made with your independent software vendor backup software.
- 4 Restart FileNet Image Services software using the method that is appropriate for your operating system. Restart stops and starts FileNet Image Services software.

Important If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Restart** button from Task Manager. (The amper-sand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Restart** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start up the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfns -y restart** at the command line or the MS-DOS command prompt.

- 5 Apply the WorkFlo queue database redo or transaction logs to roll the WorkFlo queue database forward.

Refer to your RDBMS documentation for instructions on recovering and applying redo or transaction logs.

- 6 If applying the WorkFlo queue database redo or transaction logs does not recover WorkFlo queue data up to the moment of the crash, determine if an import from optical media is desired to reinsert the document IDs and other indexes into their WorkFlo queues. Sites differ in how they prefer to recover data.

If you decide to import to recapture indexes and document IDs, use these procedures.

- a Import both the primary media and the tranlog media that have been recently written to since the backup was taken by using one of the following FileNet import utilities:
 - If you have defined WorkFlo queues for each document class and you want documents reinserted into WorkFlo queues, use the “**Import documents from media**” from the Background Job Control menu.
 - If you want more flexible options in deciding whether to insert documents into WorkFlo queues, use the **stdocimp** utility. Although the stdocimp utility imports by default to WorkFlo queues, it allows you the option to not insert documents into a WorkFlo queue if a queue is defined for the document class.
- b You do not need to identify the individual lost documents or the range of documents IDs to import, as long as you can determine which optical media should be recovered. These would be the optical media that have been written to most recently.

The import programs add missing documents to either the index database or permanent database or both, as well as to WorkFlo queues if the document class has a WorkFlo queue associated with it. The import programs do not create duplicate database entries.

Although tranlog media do not usually need to be imported, it is safer to import recently written tranlog media, in the event the disk failure occurred before the tranlog media information was completely written to the primary media. Recently written tranlog media are those written to after the backup was taken.

Tip

If you maintain a daily log of optical media surfaces which have been written to, you can easily determine which optical media to import.

Refer to the *Image Services System Administrator's Handbook* for more information on the import methods or the stdocimp utility. To download IBM FileNet documentation from the IBM support page, see [**“Accessing IBM FileNet documentation” on page 14.**](#)

WorkFlo Queue database recovery log

Use this method to restore a lost or corrupt WorkFlo queue database recovery log. It assumes your WorkFlo queue database is intact. You have two restore procedure choices – simple and complex. The simple procedure entails restoring both the WorkFlo Queue database and its recovery log in order to recover the recovery log. The more complex procedure is to only restore the WorkFlo queue database recovery log.

For the complex procedure of only restoring the WorkFlo queue database recovery log, refer to your RDBMS documentation.

The following procedure is the simple procedure. It restores both the database and its recovery log:

- 1 Stop the FileNet Image Services software by doing one of the following. You must be a member of the fnadmin or fnop and fnusr groups to stop/start Image Services software.
 - On Image Services for UNIX platforms, enter:
Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y stop** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

The Task Manager or the `initfnsw stop` command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [“killfnsw” on page 171](#), then run the `initfnsw stop` command.

Important

Rebooting your operating system may automatically start up the FileNet Image Services software. Do not assume FileNet Image Services software is stopped if you reboot the operating system. Always stop FileNet Image Services software with the Stop button in Task Manager before you back up or restore.

- 2 If your RDBMS (Oracle or Microsoft SQL Server) software needs to be restored, restore that before you restore the index database.

See your RDBMS documentation or independent software vendor product documentation for restore procedures.

- 3 Perform a restore of the WorkFlo Queue database and its recovery logs from tapes made with your independent software vendor backup software.
- 4 Restart FileNet Image Services software using the method appropriate for your platform. Restart shuts down and starts FileNet Image Services software.

Important

Start the RDBMS database first, if you have a site-controlled RDBMS database, **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Restart** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Restart** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfns -y restart** at the command line or the MS-DOS command prompt.
- 5 Apply the WorkFlo queue database redo or transaction logs to roll the WorkFlo queue database forward.

Refer to your RDBMS documentation for instructions on recovering and applying redo or transaction logs.

- 6 If applying the WorkFlo queue database redo or transaction logs does not recover WorkFlo queue data up to the moment of the crash, determine if an import from optical media is desired to reinsert the document IDs and other indexes into their WorkFlo queues. Sites differ in how they prefer to recover these items.

If you decide to recapture indexes and document IDs by importing, use these procedures:

- a Import both the primary media and the tranlog media that have been recently written to since the backup was taken by using one of the following FileNet import utilities:
 - If you have defined WorkFlo queues for each document class and you want documents reinserted into WorkFlo queues, use the **Import documents from media** from the Background Job Control menu.
 - If you want more flexible options in deciding whether to insert documents into WorkFlo queues, use the **stdocimp** utility. Although the stdocimp utility imports by default to WorkFlo queues, it allows you the option to not insert documents into a WorkFlo queue if a queue is defined for the document class.
- b You do not need to identify the individual lost documents or the range of documents IDs to import, as long as you can determine which optical media should be recovered. These would be the optical media that have been written to most recently.

The import programs add missing documents to either the index database or permanent database or both, as well as to WorkFlo

queues if the document class has a WorkFlo queue associated with it. The import programs do not create duplicate database entries.

Although tranlog media do not usually need to be imported, it is safer to import recently written tranlog media, in the event the disk failure occurred before the tranlog media information was completely written to the primary media. Recently written tranlog media are those written to after the backup was taken.

Tip If you maintain a daily log of optical media surfaces which have been written to, you can easily determine which optical media to import.

Refer to the *Image Services System Administrator's Handbook* for more information on the import methods or the stdocimp utility. To download IBM FileNet documentation from the IBM support page, see [**“Accessing IBM FileNet documentation” on page 14.**](#)

Transient database and cache – restore

Use this method to restore a lost transient database and/or cache.

These steps do the following:

- Initializes the transient database (clears all data from the transient database and its recovery logs).
- restores the transient database and cache.
- Re-enters lost documents in cache.

Initializing the transient database also clears cache. Initializing ensures that the transient database and cache are synchronized with each other before document re-entry.

You may prefer to restore the transient database and cache if the following characteristics describe your operation:

- You have a cache-only system, and
- Very little time has elapsed between the time of the last backup and the time of the disk crash, or
- You decide to rescan all lost work

If you want only to rescan batches which never fully migrated to optical media and do not want to rescan all lost work, use the alternative method, **“Transient database – initialize” on page 127**. To use the initialize method, you should have a method for determining which batches fully migrated to optical media and which batches did not.

Attention

You must have reliable backups of the transient database and cache if you are plan on initializing if a cache-only system.

- 1 Stop the FileNet Image Services software using one of the following methods. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.
 - On Image Services for UNIX platforms, enter:
Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)
 - On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.

- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y restart** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

Important

If you have a site-controlled RDBMS, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

Using the Task Manager or the **initfnsw stop** command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure **initfnsw stop** is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [**“Appendix B – FileNet system tools” on page 166**](#), then run the **initfnsw stop** command.

- 2 If you plan to restore the transient database and cache, initialize them first with the **fn_util inittrans** command. This will prepare the transient database and cache for the restore as well as prevent the recovery log from rolling forward and putting the database out of synchronization with the restored cache.

At the command line or the MS-DOS command prompt, enter:

fn_util inittrans

For a Help text description of the `fn_util inittrans` syntax, enter `fn_util` with no arguments from a command line prompt or MS-DOS command prompt.

- 3 If your approach is to restore the transient database and cache to the time of the last backup, restore the transient database and cache from tapes made with your independent software vendor backup software.

After restoring the transient database, restore disk cache immediately. The transient database and cache **must** always be backed up and restored as a pair to ensure that they are synchronized with each other.

- 4 Start the software using the method appropriate for your platform.

Important

If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.

Make sure that `TM_daemon` is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfns -y start** at the command line or the MS-DOS command prompt.

- 5 Identify all batches that have not fully migrated to optical media since the last backup and after the restore. Re-enter the lost documents.

Alternatively, you can choose the following procedure if you can't take the time to identify all batches that have already migrated to optical media:

- a Reindex and/or recommit all the restored batches.
- b Rescan, reindex, and recommit the unrestored batches

Batches failing to commit generate a "duplicate index" or "existing index" message to the event log because the index database already contains a duplicate index entry for that batch. Batches producing these error messages have most likely already migrated to optical media because the full committal process writes entries to the index database as well as to the permanent database.

However, to verify that the batches producing the error messages did commit to optical media, perform the following test:

- a Obtain the doc_id of the batch that failed during committal from the event log.
- b Retrieve and display the image associated with the doc_id.
- c If the image displays, you can delete the batch because the batch has already fully migrated to optical media and does not need to be rescanned.

- d If the image does not display but produces error messages, delete the batch and delete the document from both the permanent and index databases using the **deldocs** system tool. Then reprocess that batch again by rescanning, reindexing, and recommitting the batch because the batch never fully migrated to optical media.

See **[“deldocs” on page 167](#)** for information on the deldocs utility.

Important

If you rescan documents, do not also import the same documents from optical media. Otherwise the same document will occur twice, although with different document IDs.

Transient database – initialize

This method initializes the transient database (removes all data from the transient database and its recovery logs) and re-enters lost documents in cache. Initializing the transient database causes cache to initialize itself. Initializing ensures that the transient database and cache are synchronized with each other before rescanning.

You may prefer to use this method if:

- You lost the transient database and/or cache and have not restored the transient database or cache.
- You do not back up cache or the transient database.
- Backing up a very large page cache is not feasible in your environment.
- You have determined that it is easier for you to recover without restoring cache and the transient database.
- You have a method for determining which batches fully committed to optical media and which batches did not.

- You use optical media. Your system is not a cache-only system.

Attention

Do **not** use these steps on a cache-only system if you have no backups of the transient database and cache. **You will lose all your documents.**

- 1 Stop the FileNet Image Services software using one of the following methods. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y restart** at the command line or the MS-DOS command prompt.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Service Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

The Task Manager or the `initfnsw stop` command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [“killfnsw” on page 171](#), then run the `initfnsw stop` command.

- 2 Initialize the transient database. At the command line or MS-DOS command prompt, enter:

fn_util inittrans

For a Help text description of the `fn_util inittrans` syntax, enter `fn_util` with no arguments from a command line prompt or MS-DOS command prompt.

For information on transient database and cache restore issues, see [“Transient database and cache issues” on page 160](#).

- 3 Start the FileNet Image Services software using the method appropriate for your platform.

Important

If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start up the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfnsnw -y start** at the command line or MS-DOS command prompt.

- 4 Identify all batches that have not fully migrated to optical media. Rescan, reindex, and recommit the lost documents.

Alternatively, you can choose the following procedure if you do not take the time to identify all batches that have already migrated to optical media:

- a Reindex and/or recommit all the restored batches.
- b Rescan, reindex, and recommit the unrestored batches.

During the recommit process, batches failing to commit generate a “duplicate index” or “existing index” message to the event log because the index database already contains a duplicate index entry for that batch. Batches producing these error messages have most likely already migrated to optical media because the full committal process writes entries to the index database as well as to the permanent database.

However, to verify that these batches producing error messages did commit to optical media, perform the following test:

- a Obtain the doc_id of the batch that failed during committal from the error log event log.
- b Retrieve and display the image associated with the doc_id.
- c If the image displays, you can delete the batch because the batch had already fully migrated to optical media and did not need to be rescanned.
- d If the image does not display but produces error messages, delete the batch and delete the document from both the permanent and index databases using the **deldocs** system tool. Then reprocess that batch again by rescanning, reindexing, and recommitting the batch because the batch never fully migrated to optical media.

See [**“deldocs” on page 167**](#) for information on the deldocs utility.

Important

If you rescan data, do not also import the same data from optical media. Otherwise the same document will occur twice, although with different document IDs.

Transient database recovery log

This method is used to clear data in a lost or corrupt transient database recovery log. Clearing or initializing the recovery log puts zeros in every data block.

At startup, MKF database software determines if a rollforward of the recovery log is necessary. Rolling forward corrupt recovery log data would result in a fatal error message and Image Services would not start. If the recovery log is cleared, there is no data to roll forward to the transient database.

This method does not perform a recovery log restore and assumes:

- The transient database and cache are intact and have not been restored.
- Only the transient database recovery log is lost or corrupted.
- No current backup of the recovery log exists.

Procedures are:

- 1 Stop the FileNet Image Services software using one of the following methods. You must be a member of the fnadmin or fnop and fnusr groups to stop and start Image Services software.
 - On Image Services for UNIX platforms, enter:
Xtaskman &

Then select the **Stop** button from Task Manager. (The ampersand runs the process in the background.)
 - On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Stop** button.
 - Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y restart** at the command line or on the MS-DOS command prompt.

Important

If you have a site-controlled RDBMS database, you must stop the RDBMS software **after** you stop the FileNet Image Services software.

For FileNet-controlled databases, always use the Task Manager to stop both FileNet Image Services software and the RDBMS database. If you have a Microsoft SQL Server RDBMS, do **not** use the SQL Ser-

vice Manager to stop the RDBMS. If you do, an error displays because the SQL Service is dependent on the Image Services ControlService.

The Task Manager or the `initfnsw stop` command is the normal method to stop the system. If normal attempts fail to stop the system, evaluate your software state. Be sure `initfnsw stop` is not just running slowly; let it run for at least five minutes before determining that your system is hung. If your system is hung, run the **killfnsw** command, described in [“killfnsw” on page 171](#), then run the `initfnsw stop` command.

- 2 To clear data in or zero out the recovery log, do the following:
 - a Run the **MKF_zeroaij** utility on the transient database from a command line prompt or MS-DOS command prompt.

MKF_zeroaij zeroes out the MKF database recovery logs. Running MKF_zeroaij prevents a corrupted or lost recovery log from being applied (or rolled forward) to the transient database when the transient database next opens.

See [“MKF_zeroaij” on page 174](#) for a description and example of the MKF_zeroaij syntax.

- b If you receive a fatal error message, go to step c.

If you receive no error message, skip to step 3.
 - c If you receive the following fatal error message while attempting to run MKF_zeroaij, MKF_zeroaij has detected that the recovery log has lost data needed by the transient database to maintain data integrity:

FATAL ERROR: information in the recovery log exists
that is required for database integrity. The system

cannot allow the recovery log (formerly called the aij, or 'after image journal') to be zeroed.

To resolve the fatal error, either restore or initialize the transient database. Remember that the transient database and cache must be restored together. Refer to [“Transient database and cache – restore” on page 122](#) for procedures to restore the transient database and cache. See [“Transient database – initialize” on page 127](#) for procedures to initialize the transient database.

- 3 Start the software using the method appropriate for your platform.

Important

If you have a site-controlled RDBMS database, start the RDBMS database **before** starting FileNet Image Services software.

- On Image Services for UNIX platforms, enter:

Xtaskman &

Then select the **Start** button from Task Manager. (The ampersand runs the process in the background.)

- On Image Services for Windows Server, click on the Task Manager icon in the FileNet Image Services Server Applications folder. Select the **Start** button.

Make sure that TM_daemon is running as a service process before you invoke the Task Manager. Otherwise, Task Manager will not start up the Image Services software.

- Alternatively, for either UNIX or Windows Server, you can enter **initfnsw -y start** at the command line or at the MS-DOS command prompt.

- 4 If only the transient database recovery log was corrupted or lost (not the transient database), the transient database is now ready for production.

Multiserver restore

MultSv

A multiserver configuration typically consists of a root/index server and one or more additional servers (for example, storage library servers and application servers). The configuration can vary widely.

Stopping and restarting each server

The stopping and restart of each server when doing an offline restore in a multiserver configuration must be done in the proper sequence, as shown in the table below:

| Action | Procedure |
|---|---|
| Stop FileNet Image Services software | <ol style="list-style-type: none">1 Click the Task Manager Stop button on each non-root server in the following order: application servers followed by storage library servers.2 On the root/index server, click the Task Manager Stop button followed by the Backup button.3 On each non-root server in any order, click the Task Manager Backup button. |
| Restart FileNet Image Services software | <ol style="list-style-type: none">1 On the root server, click the Task Manager Restart button.2 Click the Task Manager Restart button on each non-root server in the following order: storage library servers followed by application servers. |

In general, you stop application servers first and restart them last. However, depending on the Image Services services running on your application server, you may need to change the sequence slightly. For example, if you run index services on an application server, restart that application server immediately after restarting the root/index server.

Procedure for a Visual WorkFlo database

Perform this additional step only if you have a multi-server system where your Oracle database contains a **Visual WorkFlo (VW)** database.

Synchronization problems may occur with an Oracle database restore in a multi-server system. It may not matter that one or all servers are restored. This is because one server may have parts of a transaction which another server does not have. Run the **VW Verify** utility after the restore to check and reconcile any discrepancies within the Visual WorkFlo database. VW Verify will find and fix more discrepancies if not all servers are restored than if all servers are restored, but VW Verify will work in both cases.

Run **VW Verify** after you have restarted the Image Services when the restore completes. On the client workstation or on the server, enter:

```
vwverify
```

To report discrepancies and commence repairs, enter:

```
vwverify -f
```

Important

if server 0 must be restored and transfer has been run since the backup, all servers should be restored. The reason is that server 0 contains work and work performer class definitions. If only server 0 is restored and the class definitions are erased, then work objects on other servers which reference these class definitions will not be usable.

If you restore only server 0, then you must run VW Verify to delete all the work objects in other VW databases which no longer have class

definitions on server 0. This deletion process may take significantly longer than restoring all servers.

For details on running the VW Verify utility, refer to the *FileNet Visual WorkFlo Installation and Administration Handbook*.

Restore issues

After your restore, your magnetic media datasets may be out of synchronization with each other or with the optical media datasets. You must ensure that your datasets are in synchronization with each other. For information on synchronizing datasets and restore issues, refer to **Chapter 5, “Synchronization and recovery information,” on page 139.**

You are responsible for your restore operation if you use independent software vendor software packages.

5

Synchronization and recovery information

This chapter contains supplemental information on the following topics:

- FileNet dataset synchronization requirements
- Synchronization problems
 - Lost data
 - Duplicate document numbers
- Updating the scalar numbers table
- Importing documents from optical media
- Comparison of FileNet and independent software vendor software coordination with MKF database rollforward recovery
- Synchronization and recovery issues
 - Permanent database issues
 - Transient database and cache issues
 - Index database recovery issues
- FileNet systems without optical storage

FileNet dataset synchronization

Dataset synchronization is defined as all datasets being updated to exactly the same point in time.

The FileNet system has multiple datasets residing on different storage media and managed by different types of software. These database storage methods are complex and interrelated. A restore of one database can introduce other problems with that same database or introduce problems to other databases. These problems include lost data or duplicate document ID numbers.

After a restore, your goals are to synchronize all MKF and non-MKF datasets and recover any data lost on your FileNet system.

FileNet dataset sequence

When a document is entered, indexed, committed and migrated to optical media on a FileNet system, the following database actions are typically performed:

- A row is created in the MKF transient database and the document is stored in cache. This transient database entry references the entire batch of documents entered.
- A row in the index database (DOCTABA table) is created. It contains the document indexes.
- An entry in a WorkFlo distributor queue is created if the system uses WorkFlo distributor queues.
- The document is written to optical media. Very often the document is written to two optical media: the primary family copy and the transaction log copy.

- The document's indexes and document number are also written to the same optical media. This is to allow a restoration of the permanent and index databases if the backup tapes are lost or destroyed.
- A row in the DOCS table of the MKF permanent database is created. It maps the document number to the two optical media locations.
- The transient database row is removed and eventually the document is removed from cache.

Synchronization requirements and examples

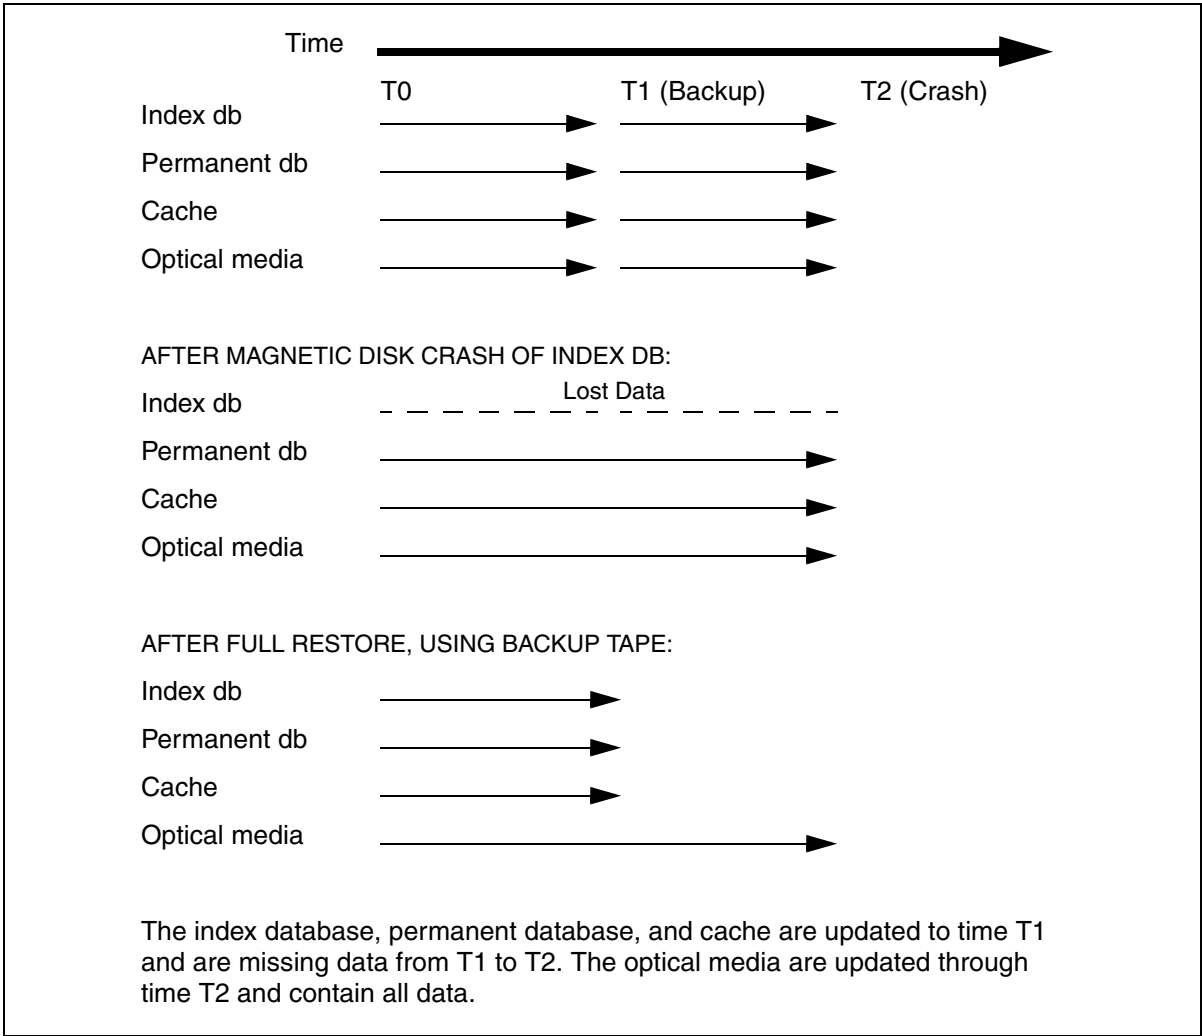
Datasets located on various magnetic disks must be synchronized or updated to the same point in time with each other. These datasets on magnetic disk must also be synchronized with the documents located on optical media.

An example of datasets falling out of synchronization is as follows: a restore is performed on the index database. That dataset is no longer synchronized with the permanent database. It does not matter if that permanent database is on the same disk as the index database. The restore sets the index database dataset contents **back** to the point in time when the backup was done while the permanent database is still at the present point in time. As a result, you must bring the index database to the point in time of the other datasets.

However, even if all magnetic disk datasets are brought into synchronization, they may still be out of synchronization with the optical media data. Optical media contain all documents, including those migrated after the backup was performed. Thus, the documents on optical media can be more current than the restored permanent and/or index datasets. Therefore magnetic disk datasets fall out of synchronization with the optical media data after a restore.

See **“Representation of Unsynchronized Datasets” on page 143** for an illustration on how datasets fall out of synchronization. In this graphic, the index database on magnetic disk has experienced a disk crash at the T2 point in time. The disk crash has caused the index database to lose all data. A backup tape was taken at end of T1 point in time. A full restore is performed from the backup tape so that all datasets are synchronized to point T1 except for the optical media.

Now, the index database, the permanent database, and cache are out of synchronization with the optical media. The optical media were not affected by the crash. They were updated through time T2. The three datasets were restored to time T1.



Representation of Unsynchronized Datasets

Synchronization problems

When your datasets on magnetic disk and optical media are not synchronized after a restore, two problems may be encountered:

- Lost data since the last backup.
- Potential duplication of document ID and optical media surface ID numbers during subsequent processing.

Lost data

After a restore, you have lost data that was processed since the last backup. The restored datasets will not be synchronized with any other intact magnetic disk datasets that were not restored. The restored datasets will also be out of synchronization with the optical media data. Even if you do a daily backup, documents have probably been processed and written to optical media in the hours since your last backup. These recent documents will be lost or irretrievable to the FileNet system because the permanent and/or index database entries needed to retrieve them were lost in the restore. Refer to [Chapter 4, “Advanced restore methods,” on page 83](#) for specialized ways to recover lost data.

Duplicate document numbers

You run the risk of your FileNet system generating duplicate document IDs and duplicate optical media surface ID numbers in any restore scenario involving the permanent database.

Background information on ID numbers

Whenever a new document is entered into the system, FileNet Image Services assigns it a document ID number. This document ID number

originates from the scalar numbers table (SNT) in the permanent database. The SNT also assigns new surface ID numbers as well as the background request numbers (also called background job ID number). These numbers are incremented with each new issuance. This ensures that the value in the SNT is always greater than the last document ID and surface ID.

If a disk crash occurs that results in the loss of the permanent database, the SNT is also lost. When the permanent database is restored from the last available backup tape, the old SNT is also restored. This means the next available document ID and surface ID in the restored SNT have been previously assigned and you have the capability of getting duplicate IDs assigned. **The values in the SNT must be greater than the last assigned document ID and surface ID.** Any rolling forward of the transaction log won't necessarily roll forward to the exact moment of the crash.

To maintain the integrity of your data, you must not resume document entry until this problem is corrected.

Tip

You can use the MKF_tool utility to display the current contents of the scalar numbers table. Following is an example command to display the entire contents:

```
MKF_tool> select scalar_numbers *
```

snt.chkpt file (SNT checkpoint)

To guard against outdated SNT values, FileNet Image Services updates a backup file as it updates the SNT. This backup file is known as the scalar number checkpoint file or snt.chkpt. This checkpoint file is located in /fnsw/local/sd/snt.chkpt (or \fnsw_loc\sd\snt.chkpt on Win-

dows Server). The file is timestamped to help you determine how old the checkpoint file is compared to the last restore of the SNT.

The snt.chkpt file values include the next available document ID, next optical media surface ID, next background request number, and a number representing the date/time of last update to the snt.chkpt file, as shown in the example below.

```
5018595 3344 98 856828856
```

Each time FileNet Image Services starts, Document Services automatically compares the values in the snt.chkpt file to the values in the SNT of the permanent database.

If any of the snt.chkpt file values is higher than those in the SNT, Document Services terminates and logs the following message to the system log or event log:

```
Severe Error condition: The Scalar Numbers Table is behind the snt.chkpt
file. This should only happen after a Permanent DB restore has been done.
Continuing with this condition may cause multiple documents to be
committed with the same doc ID. To resolve this problem, you must update
the Scalar Numbers Table with the SNT_update program. Doc Services will
not function until this problem is resolved.
```

Updating the scalar numbers table

You need to update or advance the document ID number in the scalar numbers table if you have restored the permanent database or performed an entire disk restore or a full restore. This is accomplished with the SNT_update command.

Running SNT_update

The SNT_update utility updates the contents of the SNT using values from the snt.chkpt file. If the snt.chkpt file survives a magnetic disk crash and the permanent database does not, you can use the snt.chkpt file to restore the SNT to its current state.

After running the SNT_update utility, the system will not assign duplicate document ID and surface ID numbers. You may proceed with document entry of any lost data. After these documents have been rescanned, reindexed, and recommitted, datasets on magnetic disk and optical media will be synchronized.

After you run the SNT_update program, you see the following example display. The values for surface, job, and document IDs are for illustration only. The values in your display will be different:

```
corona>SNT_update
"The time stamp on your checkpoint is Wed May 23 13:45:52 2001
"The SCALAR_NUMBERS table values in the Permanent Database are:
      Next Surface  ID (MKF) = 3344
      Next Job      ID (MKF) = 98
      Next Document ID (MKF) = 520500
The check point file values are:
      Next Surface  ID (checkpoint) = 3344
      Next Job      ID (checkpoint) = 98
      Next Document ID (checkpoint) = 530500

      ===> New Next Document ID will be = 531500

Do you wish to update the Permanent Database Scalar numbers with the
checkpoint
values? (y/n):
```

The time stamp displays the date and time of the last update to the `snt.chkpt` file. The time stamp can help you determine how old the `snt.chkpt` file is compared to the values in the restored SNT.

The next document ID number displayed in the above example will be the checkpoint value of 530500 plus 1000 (or 531500). If the checkpoint document ID value is lower than the MKF value by 1000, then `SNT_update` retains the MKF document ID value. In this example, the surface ID and job ID numbers did not require updating. If the surface ID and job ID checkpoint values were higher, they would replace the MKF values.

In the example illustrated, after you enter **y** to update the SNT (MKF) values with the `snt.chkpt` (checkpoint) values, the following confirmation displays:

Scalar numbers table updated

After the update occurs, you may restart the FileNet Image Services software to restart document services.

If you enter **n**, `SNT_update` terminates.

If you enter **y under other circumstances**, any one of the following messages may display indicating successful completion or errors that prevent the successful update of the scalar numbers table:

- “Scalar numbers table updated”
- “Update was not necessary
Scalar numbers table already up to date”
- “Scalar numbers table update failed, err=<err_code>”

where <err_code> is a value indicating the type of error encountered. SNT_update terminates after displaying this message.

Updating the snt.chkpt file

A full system restore overwrites the current values in snt.chkpt with old values. The same would occur if the /fnsf file system (or \fnsf_loc) was included in a full disk restore. **Do not run SNT_update if the /fnsf/local (\fnsf_loc) partition has been lost or damaged and then subsequently restored.** FileNet Image Services will try to update the SNT with old, recently restored values.

In this scenario, you have to manually edit the snt.chkpt file. The procedures for updating the snt.chkpt file are found in the appropriate restore method section (Simple vs. Advanced).

Importing documents from optical media

Importing documents from optical media is the preferred method of recovering lost data under the following conditions:

- To prevent redundancy of the same documents on optical media.
- To restore a large volume of data processed since your last backup.
- Your backup tapes are lost, destroyed or unreadable. You can rebuild the index and permanent database from optical media.

Procedures for importing documents from optical media and information on the import utilities are found in the appropriate section in **[“Restore methods by dataset type” on page 89.](#)**

Import issues

Importing from optical media recovers most, if not all, data. You should be aware of the following issues:

- You lose any changes to indexing information made after migration to optical media has completed.

If you do not index your documents immediately before executing committal and you restored the index database, you must import **first** from optical media those documents that were committed after the last backup. Then reindex those documents that were imported because the index values on optical media will not be current.

- You lose all information on folders created since the last backup.

Because no folder information resides on optical media, you lose all folder data, for example, inserts, deletions, and renames. You must manually recreate the data or write a program to recreate it.

- If any unmirrored cache or transient database was on the disk that was replaced, you must rescan any uncommitted batches because documents in uncommitted batches have not been written to optical media.
- If you used the **odrecover** utility to import from optical media, your related data was not placed into WorkFlo distributor queues.

If you want the document IDs inserted into the WorkFlo queues, re-import the optical media using an import utility such as **stdocimp**.

- Any optical media put into the storage library after the date of the restored backup tape must be incorporated prior to import. Otherwise, the system won't recognize the relatively new platter. Incorporation identifies the surface to the system and allows the surface to be imported. To incorporate optical media, use the Server Con-

trol Program (SCP) for FileNet Image Services releases prior to 3.3.0 or use the Storage Library Control (SLC) utility for FileNet Image Services releases 3.3.0 and later. Both utilities are documented in the *IMS System Administrator's Handbook* for your FileNet Image Services release.

Import assumptions

The assumptions regarding your system are:

- You restored the permanent database or the index database or both.

After the restore, your databases are not synchronized with each other or with the optical media. WorkFlo queues reside either in the index database or on a WorkFlo queue server.

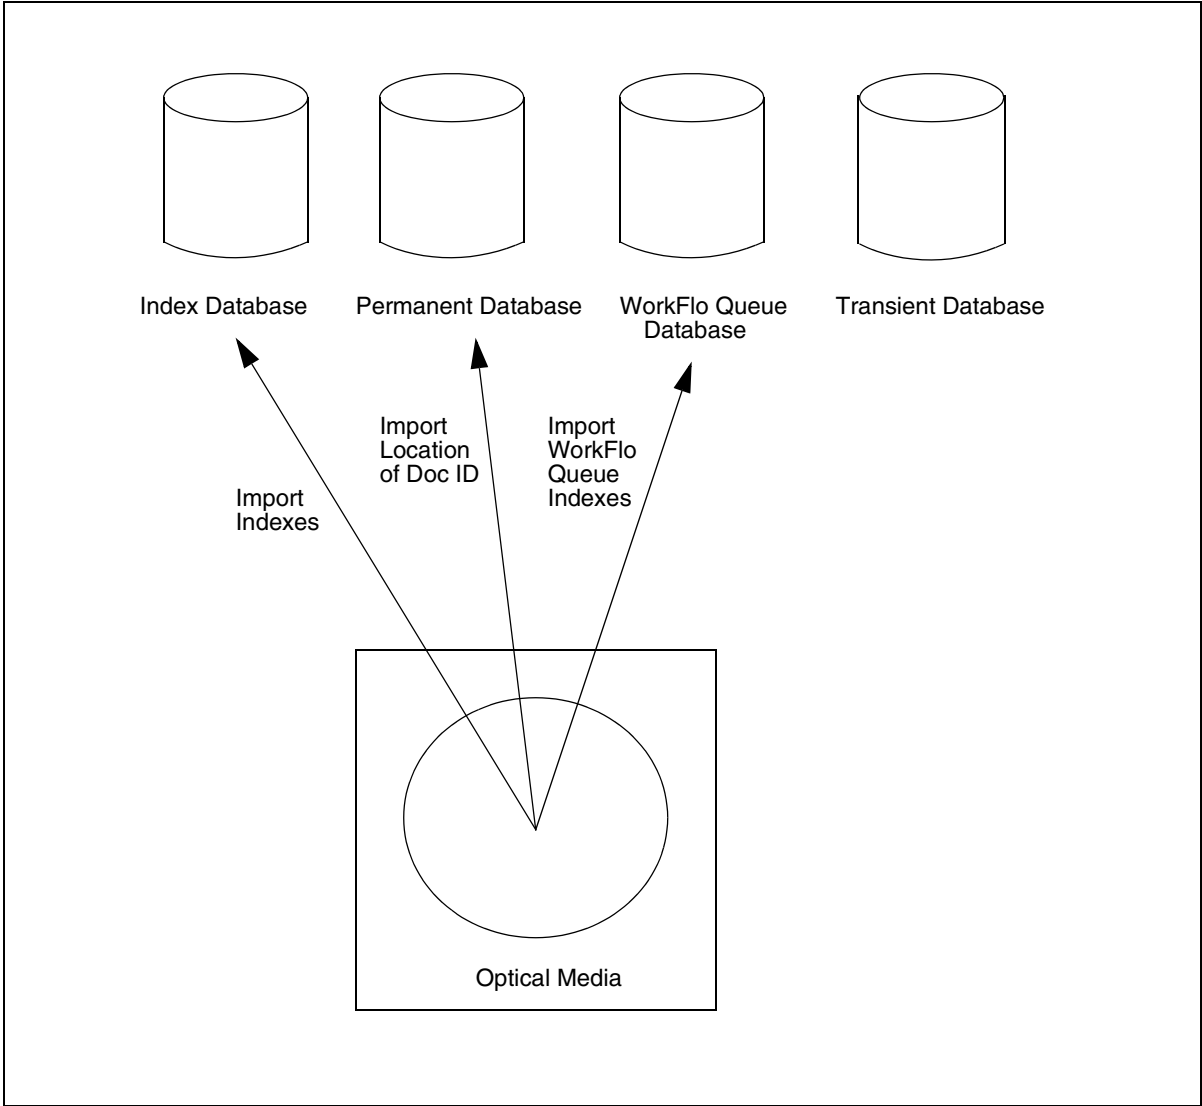
- You have lost document references in the permanent and/or index database. You want to restore them by importing from optical media.
- Your system is a stand-alone system.

A stand-alone system has not been set up as a compatible system with other systems. No system serial numbers (SSN) are in the file, `/fnsw/local/sd/as_conf.g`, in the “compatible systems” section on the storage library server.

Import diagram

The diagram, **“Representation of Importing Process” on page 153**, shows what type of data is imported to the index database and permanent database. If the document class for the optical media has WorkFlo distributor queues in the document classes definition, the import process also imports the WorkFlo queue indexes. The transient

database is not affected. No data or images are inserted to cache or the transient database during an import from optical media operation.



Representation of Importing Process

MKF database rollforward recovery

You should understand how a lack of coordination between MKF dataset, MKF recovery log, and restore timing can contribute to lost data and dataset synchronization problems.

This section describes:

- The general rollforward process.
- FileNet and independent software vendor software coordination issues with MKF rollforward recovery.
- Synchronization issues after a rollforward.
- Specific restore issues that apply to an MKF database.

MKF rollforward process

The permanent, security and transient databases are MKF (multi-key file) databases. Each MKF database has an MKF recovery log. Most systems are configured with one recovery log. However, you may have more than one.

The next time an MKF database is opened after a restore, FileNet Image Services detects that the MKF recovery log is more recent than the database. MKF then attempts to automatically roll the database forward from the recovery log. This process is also known as “splicing in” the MKF recovery logs to the MKF databases. MKF attempts to recover data up to the end of the last transaction. All committed changes are rolled forward and automatically applied to the database.

After the MKF database rollforward completes, your MKF databases are synchronized with each other and no data is lost from MKF databases.

Rollforward may not be possible if the MKF recovery log is on a disk that crashes or if the recovery log is overwritten during normal system activity.

FileNet MKF coordination

FileNet Image Services coordinates with FileNet Enterprise Backup/Restore software to warn the user when a potential for losing MKF recovery log data is detected.

Changes to an MKF database are first written to its corresponding MKF recovery log. The MKF recovery log writes to (or flushes) to the MKF database at the following events:

- a Enterprise Backup/Restore is executed.
- b FileNet Image Services is brought down.

An MKF recovery log contains database updates since the last time FileNet Image Services software was recycled or since the last backup using FileNet Image Services backup tools. This allows for a restore from the last backup and roll forward from the recovery log.

Attention

The recovery logs will not know that you are using a third party backup program. If you do your third party backups on-line, the MKF recovery logs will **NOT** write to the MKF databases unless you bring down FileNet Image Services.

When using third party backup software, you should do **off-line** backups so that the recovery logs will write to the databases.

An MKF recovery log has a finite size. When the data in a recovery log exceeds the log's finite size, **the log will overwrite itself**. This is

known as the **trigger point**. It is the point at which the initial recovery log records required for roll forward from the most recent backup are overwritten. This occurrence should be avoided. Third party backup programs put your system at greater risk of reaching the trigger point.

Tip

You do have the capability of adding additional recovery logs.

When the recovery log is overwritten, a roll forward after a restore is not possible. When the trigger point is reached, an informational message is written to the system log to warn you that you cannot roll the database forward. This message is your warning to immediately bring down FileNet Image Services (to flush the recovery log data to the database) and make an immediate backup of the database (not the recovery log) in order to preserve the database changes now lost in the recovery log.

The trigger point is updated when you bring down FileNet Image Services or perform a backup using Enterprise Backup/Restore. You can safely continue normal FileNet Image Services operations until you reach the end of the recovery log, which is the next trigger point.

Recommendations to prevent overwriting are covered under **“MKF recovery log requirements” on page 158.**

MKF summary

The following points summarize the coordination between FileNet backup/restore software and MKF database software:

- The FileNet backup software informs MKF when backup begins and ends. MKF cooperates with the FileNet backup software to make on-line backups and interval backups possible. **This does not apply to third party backup software.**

- During each FileNet backup operation, MKF computes and saves a trigger point in the database control block. MKF uses the trigger point to determine whether the recovery log data is current or not, and whether it can be rolled forward to the end of the last transaction. **This does not apply to third party backup software.**
- An informational message is written to the system log if, during normal updates to the MKF databases, enough recovery log data is produced to reach the trigger point.

This message informs you that you can no longer roll the database forward by using the disk-resident recovery log after a restore. This message appears only once, no matter how many times the recovery log wraps around. You won't see the message again until you perform a backup (with Enterprise Backup/Restore) and enter another overwrite condition.

This message is a warning to you to do an **immediate backup** to avoid losing any data. It is also an indication to increase your backup frequency and/or increase the size of the recovery log. You will not normally see the warning message if you bring down your FileNet Image Services software, do frequent backups (with Enterprise Backup/Restore) and the size of your recovery log is adequate to hold at least two or three days' worth of processing.

- When MKF opens an MKF database after a restore, a message is written to the event log alerting you whether or not the recovery log spliced in during the startup. You are informed whether or not the rollforward successfully occurred.

MKF recovery log requirements

You absolutely must have measures in place to prevent the overwriting of the MKF recovery logs, especially if you are using a third party backup product.

To improve the chances of an MKF roll forward recovery, observe the following recovery log requirements:

- Ensure that the MKF recovery logs reside on a different physical volume than the MKF database.

The recovery logs contain copies of all changes made to the database itself. If you lose the database, but not the volume that contains the recovery logs, you may be able to recover up to the moment of the crash.

- Allocate enough disk space to MKF recovery logs to hold at least two to three days' worth of recovery log data.

The recovery log must be large enough to hold all of the changes that occur between two events when you bring down FileNet Image Services. If the recovery log is not large enough, some changes will be overwritten and lost after a restore, and roll forward after a restore will not be possible. The MKF recovery log is not automatically archived. When the log is full, recording wraps and overwrites the contents of that log. If you have two recovery logs, recording continues on the second log. When the second log is full, recording wraps to the first log and overwrites the contents of the first log.

Independent software vendor product coordination with MKF

MKF database software is designed to work with FileNet Enterprise Backup/Restore software. MKF recovery logs won't work the same

way with independent software vendor backup/restore software as they do with FileNet backup/restore software. A independent software vendor backup/restore package is not able to communicate and coordinate effectively with MKF to warn you if you are overwriting recovery log data.

The major areas where independent software vendor backup/restore software may not coordinate with MKF are the following.

- The independent software vendor backup/restore software does not inform MKF when a backup has occurred. The result is that MKF will **not reset the trigger point** in the MKF recovery log.
- MKF writes a **message** to the event log the first time the trigger point is encountered. This is the point when the recovery log begins to be overwritten. **No other messages** will be sent because the trigger point is not reset.

The message may be ambiguous because, at this point, you do not know whether you can roll forward, since you do not know if you completed a backup before the trigger point was reached.

The warning message is intended to notify you that your recovery log is not large enough or that you have gone too long without a backup.

- MKF does not know and is not able to notify you whether or not an MKF rollforward recovery has taken place.

Precautions against Data Loss

- Users of independent software vendor backup/restore packages should take the following steps to decrease the chance of losing data in the event of a system failure:
- Do full **daily** offline backups.

- If you do daily backups, make sure your MKF recovery log is **large enough** to hold at least **two to three days** worth of data. This protects the data in case your last backup is not good. Other than the initial warning, you will not be warned if your recovery log is too small and is overwritten.

Synchronization and recovery issues

This section discusses some of the synchronization and recovery issues that result if you restore the permanent database, transient database, and cache and apply transaction logs to the index database.

Permanent database issues

Unless you can determine whether your permanent database has rolled forward, your permanent database may be out of synchronization with your index database. When using independent software vendor backup software, you probably will not get a message in the system log informing you whether or not your permanent database has been rolled forward.

To synchronize the permanent database with the index database, you can apply Oracle redo logs or Microsoft SQL Server transaction logs to your index database. For Oracle databases, you can use Oracle's roll-forward utility to roll the index database forward.

Transient database and cache issues

The following methods ensure that your transient database and disk cache are synchronized and you recover lost data:

- Back up and restore the transient database and cache together.

- Reinitialize the transient database and cache in lieu of some backup/restore situations.

Reinitializing transient database and cache

You may choose not to back up and restore the transient database and cache if you have a very large page cache; performing a daily backup of these may not make good business sense. Under that circumstance, if you lose the transient database and/or cache, you can recover them as follows:

- 1 Reinitialize the transient database and its recovery logs to zero. This ensures that the transient database and cache are synchronized with each other before rescanning. Reinitialize the transient database with the system tool, `fn_util inittrans`.
- 2 Determine which lost documents need to be rescanned.

To simplify determining which batches did not get written to the index database, you can commit the batch ID as an index with the document. Committing the batch ID as an index helps you determine which batches were not written to the index database.

- 3 Re-enter the documents lost from cache.

Typically, so much work gets processed after the backups of the transient database and cache that restoring them would cause duplicate documents to be recommitted to the optical media. Also you do not need to back up and restore the transient database and cache if you intend to re initialize them after the restore. Refer to **“Transient database – initialize” on page 127** for detailed procedures.

Restoring transient database

If you restore the transient database, your transient database will be out of synchronization with your disk cache unless you **prevent** the transient database from rolling forward. You prevent the roll forward by initializing the transient database and its recovery log with the **fn_util inittrans** system utility. By initializing the recovery log, you do not have to worry about the transient database rolling forward and putting the database out of synchronization with the restored cache. Refer to [**“Transient database and cache – restore” on page 122**](#) for detailed procedures.

The transient database must not be rolled forward because cache will not roll forward. FileNet Image Services tries to roll the transient database forward on the next database open after a restore and may cause the transient database to be unsynchronized with the disk cache.

Restoring disk cache

Although you can synchronize the permanent and security databases by rolling them forward, you cannot do so for disk cache. Cache **cannot** be rolled forward. Your primary protection against disk crashes must be disk mirroring, so you should set up disk mirroring on your entire FileNet system.

Index database recovery issues

If you enabled archiving of Oracle redo logs, you can apply the logs to your index database to attempt to synchronize your index database with the MKF databases, in particular, the permanent database. Rolling the index database forward via Oracle's roll forward utility should synchronize the index database with rolled forward MKF databases.

Oracle control files must be handled in the following manner to ensure that the appropriate, current control file is in place after a restore:

- As a regular safety practice, mirror the Oracle control file and put each copy on a separate disk.
- After a crash, recover the control file from the good copy that you duplicated, if possible, instead of restoring the bad copy.
- Refer to Oracle's Server Administrator documentation for more information on Oracle control files.

All Microsoft SQL Server databases and files must be restored as described in the table, **"Microsoft SQL Server Databases to Back Up" on page 63**. Refer to the *Microsoft SQL Server Administrator's Companion* for more information.

FileNet systems without optical media

If your FileNet Image Services system does not use optical media, then all work in progress, including committed documents, are stored in magnetic disk cache. Systems without optical media need to synchronize the location of cache documents with the document information in the RDBMS and the MKF databases.

Systems without optical media should observe the following guidelines to synchronize, protect, and recover data:

- You **must** perform full, offline, daily backups of all data, including disk cache and transient database.
- You should mirror cache.
- You must rescan and reenter any data lost after a restore. There is no optical media data to import.

Appendix A – Support policy for backup/restore

The FileNet Enterprise Backup/Restore (EBR) program is provided for backing up and restoring customer data. EBR provides incremental backup, online backup, and backup to a file.

Important

If you are upgrading to the current FileNet Image Services release, you need to ensure that your archived backups are in EBR format.

You also have the option to use a independent software vendor product to back up and restore your systems. If you use a independent software vendor product, you take full responsibility for the protection and integrity of your data.

Precautions

If you plan to use a independent software vendor product, we recommend the following three precautions:

- 1 Perform occasional full backups using EBR.
- 2 Follow the complete test procedure, outlined in **Chapter 1, “Requirements and recommendations,” on page 16** of this manual, before switching to a independent software vendor product on a production basis.
- 3 Retain the FileNet backup tapes, from Step 1, in a safe storage area. The archived tapes should be kept reasonably current (no more than

six months old). The archived tapes will provide a baseline in case you encounter problems in restoring your system.

Who to contact

There are a number of independent software vendor backup and restore products on the market. Since IBM does not possess the in-depth technical skills to diagnose Third Party software problems, we must refer our customers to those Third Party software vendors for technical support when the problem is not with IBM software.

Recommendation

To minimize potential problems, you should use the Simple Backup and Restore approach outlined in **Chapter 2, “Simple Backup and Restore,” on page 27** of this manual. Full support is available for the FileNet backup and restore utilities and EBR.

Appendix B – FileNet system tools

This appendix contains descriptions of the FileNet Image Services system tools referred to in this manual. These tools are used for analysis, diagnosis, update, and repair of FileNet Image Services components. You may need these tools to complete your backup, to restore operation, or to synchronize your datasets after a restore.

Important

These tools should be used by a trained system administrator or support representative. Incorrect use can damage your system.

The tools are listed in alphabetical order. The descriptions provide syntax, usage information, sample output and/or examples, precautions, and other information as appropriate.

The following table summarizes the tools in this appendix.

FileNet System Tools

| Tool | Description |
|-------------|---|
| deldocs | Deletes documents from the index database and permanent database. |
| killfns | Forcefully shuts down the FileNet system. |
| MKF_zeroaij | Clears any data that exists in either the transient or permanent after-image recovery logs. |
| odrecover | Imports documents from optical media for use in rebuilding damaged or corrupted permanent or index databases. |

deldocs

Description

The deldocs tool deletes documents from the FileNet Image Services index database and/or the MKF permanent database. It may also delete selected documents from cache, given selected settings.

Use

Use deldocs to delete documents that have been committed using the “no cataloging” option of Database Maintenance.

You can use deldocs to delete a single document, multiple documents, or all documents for a specified media surface. To delete multiple documents, specify a list of document IDs in a file that deldocs uses as input.

You can also delete all documents on a media surface and use Background Job Control Erase Media or Consolidate Media to erase the surface (only when using erasable media). The Consolidate Media function provides an option to erase the media as it consolidates the media.

Attention

This command deletes all references to the documents from the FileNet Image Services index and MKF permanent databases. You can, however, re-import these deleted documents from storage media.

For details on deleting documents from cache, see [“Documents in cache” on page 168.](#)

Syntax

deldocs [-f <filename>] [-s <surfaceid> [-b]]

- f <filename>** Text file name containing document IDs, one ID per line
- s <surfaceid>** Delete all documents on the surface specified by <surfaceid>. This option does not erase documents from the media surface.
- b** Delete documents from both sides of the surface specified by <surfaceid>. Must be used in conjunction with the -s <surfaceid> parameter.

Documents in cache

The deldocs -f command behaves differently depending on two things:

- 1** Whether the class assigned to the document is set to migrate to optical disk, and
- 2** Whether the operator overrides this document class setting by modifying the Capture Committal Component.

Document class setting

When scanned into the FileNet system, the document is assigned to a class. The document then inherits any parameter values set for the class to which it is assigned.

The System Administrator configures classes using the Database Maintenance application, available through the Application Executive (Xapex). The Database Maintenance Classes option, Define/Update Document Classes, enables the administrator to set the Migration to O.D. field to either Yes or No:

- Yes specifies that any documents associated with this class will be migrated to optical disk. It also flags the document as "ageable," unlocking it in cache. As long as the Capture application does not override this setting before committal, deldocs does not delete the document from cache. Instead, either CSM_daemon can remove the document from cache at an appropriate time or you may remove the document manually using the CSM_tool.
- No specifies that any documents associated with this class will not be migrated to optical disk. After committal has completed successfully, the document resides in cache as a "locked" document. As long as the Capture Professional application does not override this setting before committal, the deldocs command deletes a locked document from cache.

Tip

Use the CSM_tool utility to see if a document in page cache contains a "locked" or "ageable" status.

Capture committal component

Under normal conditions, the Capture Committal Component inherits the default setting for the given document class, whether Migration To O.D. is set to Yes or No. If set to Yes, Capture shows Commit with Migration checked. If set to No, Capture shows Commit without Migration checked. If the operator modifies the Capture Committal Component, it overrides the setting assigned to the document class.

Checklist

Before you use deldocs, be aware of the following:

- You must be logged on to FileNet security (Windows Server users, logon through the Application Executive) before using deldocs.

- Even if a document is not in both the index and document databases, deldocs deletes the document.

Procedure

- 1 Log on with appropriate user privileges.
- 2 Enter the deldocs command with desired options.

killfnsw

Description

The killfnsw tool is issued at the system command line to forcefully stop components of the FileNet system such as processes, daemons, databases, and users.

Attention

Do not use killfnsw unless you are absolutely sure normal shutdown methods have failed. We recommend that killfnsw be used only by an experienced system administrator or support representative to avoid damage to your FileNet system.

The killfnsw program removes all existing interprocess communication structures including shared memory segments, message queues, and semaphore sets used by the FileNet Image Services software. It abruptly shuts down databases.

If you run killfnsw with no arguments, the Task Manager daemon (TM_daemon) remains active. You must use a killfnsw argument to terminate TM_daemon.

Use

Use killfnsw **only** after normal attempts to stop the software fail and you have evaluated the reason your system cannot stop normally. For example, make sure the normal shutdown process or command is not just running slowly. (Allow at least five minutes for a shutdown process or command to complete before resorting to killfnsw.) Use ipc_tool to analyze your system's state. However, if ipc_tool does not run or cannot complete, you may be forced to use killfnsw to stop your system or components of your system.

You may be forced to use killfnsw to terminate the TM_daemon. Normally, the TM_daemon continues to run across shutdown cycles of FileNet Image Services software. On occasion, you may have to terminate TM_daemon as well. For example, when you need to install fix packs that include shared libraries on the FileNet system, you must terminate all processes that use the shared library before installing the new shared library. Since the TM_daemon uses some of the FileNet shared libraries, you must terminate it to upgrade the shared library or install the new shared libraries.

Tip For more information about starting and stopping a service process in a Windows Server environment, see the *System Administrator's Companion for Windows Server*. To download IBM FileNet documentation from the IBM support page, see [**“Accessing IBM FileNet documentation” on page 14.**](#)

Syntax

```
killfnsw [-c] [-D] [-v] [-S] [-A] [-p <pid>] [-y]
```

- c** Kills only clean shared libraries with zero use counts (AIX/6000 platforms only)
- D** Kills TM_daemon
- v** Runs in verbose mode
- S** Kills processes when the FileNet system is in single user mode
- A** Removes all interprocess communication (IPC) segments
- p<pid>** Kills the process specified by process ID <pid>
- y** Suppresses confirmation messages

Checklist

Before you use killfnsw, be aware of the following:

- You should use killfnsw **only** after normal attempts to stop the software fail and you have evaluated the reason your system cannot stop normally.
- In a UNIX environment, running the whatsup command restarts a TM_daemon that you terminated with killfnsw.

Procedure

- 1 Use Task Manager to list all running FileNet processes.
- 2 Terminate FileNet processes.

Use Task Manager or the **initfnsw stop** command.

- 3 Run killfnsw with appropriate options to terminate active processes.

When killfnsw completes, run whatsup to determine if the processes have been stopped.

Important

If you run killfnsw with no arguments, the Task Manager daemon remains active.

- 4 Terminate the Task Manager daemon, if necessary, enter the following command:

killfnsw -D

Important

Do not run the whatsup command after you terminate TM_daemon in a UNIX environment. Doing so automatically restarts TM_daemon.

MKF_zeroaij

Description

MKF_zeroaij exists for one reason: to zero the after-image journal (aij) files after a disk spindle on which the aij resides has failed. The aij is automatically zeroed by FileNet processes for any other situations that require the aij to be zeroed.

Tip “Recovery log” refers to the after-image journal.

When you start MKF_zeroaij, the MKF_zeroaij program zeroes the after-image recovery log. The program requests the base data file name rather than the recovery log file name because the base data file has the names of **all** recovery log files and data files. The base data file name is in the ddl file for the database (for example, /fnsw/local/sd/1/permanent.ddl or /fnsw/local/sd/1/transient.ddl).

Use

The MKF_zeroaij program clears any data that exists in either the transient or permanent recovery logs. If the magnetic media containing the recovery log fails, a new drive must be initialized prior to running normal operations. Use MKF_zeroaij to initialize the new disk drive.

Syntax

MKF_zeroaij <base_data_file>

<base_data_file> Identifies the full path name of the base data file of the MKF database. Be sure to specify the base data file name and not the after-image journal file name.

Example

MKF_zeroaij /fnsw/dev/1/permanent_db0

Sample output

```
>MKF_zeroaij /fnsw/dev/1/permanent_db0  
zeroing /fnsw/dev/1/permanent_rl0 (block 0 through 40959)  
MKF_zeroaij: done
```

Checklist

Before you use MKF_zeroaij, be aware of the following:

- You must notify the FileNet Response Center before you attempt to run MKF_zeroaij.
- You must use the base data file name, **not** the recovery log file name.
- You must run MKF_zeroaij at the server on which the database resides.
- You must stop the database before running this command.
- No processes should be linked to MKF when MKF_zeroaij is started. Normally, you should stop FileNet Image Services software before you start MKF_zeroaij.

Procedure

- 1 Call the FileNet Response Center.
- 2 Shutdown the FileNet Image Services software.
- 3 Run MKF_zeroaij.
- 4 Restart the FileNet Image Services software.

odrecover

Description

The odrecover tool initiates a docimport background job to start importing documents from storage media. This works the same as calling the stdocimp tool with selected options. It takes the following actions:

- Updates the next available document ID in the scalar_numbers table, if necessary.

If you have to restore your system from backup tapes, the current document ID on the backup tapes usually does not match the document ID that was current when the database was damaged.

- Avoids importing deleted documents when only one of the document tables (DOCS or DOCTABA) is damaged.
- Verifies that a document is in both document tables.

If the document is in only one table, odrecover adds it to the other table. If the document is not in either table, odrecover does not add the document to either table. If both document tables are damaged, use the `-d` flag to import all documents on the storage media.

- Prevents adding documents to WorkFlo queues.

Re-imported documents should already be in the appropriate WorkFlo queues (or already processed by WorkFlo).

You can have odrecover verify that the associated document classes in the importing and exporting systems match during import. By default, odrecover does **not** require consistent document classes.

Important

If odrecover encounters an error (such as mismatched document class) during the import of a document into DOCTABA, odrecover still corrects an inconsistency in the databases based on the DOCTABA table in the following manner. If the document exists in DOCTABA but not in DOCS prior to the import, the document is imported to the DOCS table. If the document exists in DOCS but not in DOCTABA prior to the import, odrecover **deletes** the document from the DOCS table. Check for message tuple 90,0,18 in the following logs:

- system event log
- import log file
/fnsf/local/logs/bkglog/ImpLog.xxxxxx for UNIX platforms
<drive>:\fnsf_loc\logs\bkglog\ImpLog.xxxxxx for Windows Server platform where xxxxxx is the odrecover job number.

To recover the deleted documents, run odrecover -d.

If you run odrecover with the exact match flag (-e), odrecover issues message 90,0,18 when a mismatch exists between the index entries of the importing and exporting systems. The program logs messages in both the system error log and in the Background Job Control import log (ImpLog.xxxxxx) for the odrecover job.

Samples from both these logs are shown below:

```
97/02/12 16:23:10.000 90,0,18 <fnswg> docimport 111 (0xa4.84) ...
Docimport Job#=111, surf=3344; doc 1167956, ssn 8008, Can't get index
description, name='stringA'

97/02/12 16:23:10.000 90,0,18 <fnswg> docimport 111 (0xa4.84) ...
Docimport Job#=111, surf=3344: Translate descriptor to DIR failed.
Document id=1167956, ssn=8008

97/02/12 16:23:10.000 90,0,18 <fnswg> docimport 111 (0xa4.84) ...
Docimport Job#=111: Import: DB insert failed. Document id=1167956,
ssn=8008

97/02/12 16:23:46.000 76,0,262 <fnswg> docimport 111 (0xa4.84) ...
Document import Job 111 completed
```

The following entries are from the Background Job Control import log file:

```
Information for import job number 111 started at Wed
May 9 16:22:40 2001

Job parameters:
Input surface id: 3344
Import from both sides: no
Import all files: yes
Document class must match exactly: yes
Reset surface statistics: no
Insert into doctaba: yes
Security options: from document class

(continued on next page)
```

(continued from previous page)

```
Update scalar numbers tables: yes
Update WorkFlo queue: no
Skip import of deleted documents: yeys
Number of documents to skip on media: 0
High priority: no

Errors encountered:
document id: 1167956, ssn: 8008, error: <90,0,18>

End of job information:
Number of errors encountered: 1
Number of documents read from input media: 610
Number of documents imported: 0
Number of duplicate documents in docs table: 1
Job completed at Wed May 9 16:23:46 2001
```

Use

Use `odrecover` to import documents from media and rebuild the permanent and index databases when either or both databases have become damaged or corrupted.

Syntax

```
odrecover [-b] [-d] [-e] <surface ID>
```

- b** Imports both sides of the specified medium
- d** Imports all documents on the specified media, regardless of whether the documents are in the DOCS or DOCTABA table. Imports all documents, even those that have been deleted from the databases. This flag is usually specified if both document tables are corrupted.

- e Verifies that the importing and exporting systems have an exact document class match for each document to be imported. If mismatch occurs, odrecover logs an error message and the document is imported or deleted based on the DOCTABA table. If you do not specify this option, odrecover continues with the operation even if the document classes do not match.

<surface ID> Specifies the media surface to import (for example, 3002)

Checklist

Before using odrecover, be aware of the following:

- Entering odrecover with no flags displays online help that includes a description of odrecover and a definition of each parameter.
- When you run odrecover, you must import documents from all media that were written since the last backup was made. If you do not import all media, the assignment of duplicate document IDs may result.
- After running odrecover, re-scan documents that were in page cache (that is, those documents not yet written to media).
- The odrecover tool does not import documents into WorkFlo queues.
- The default behavior of odrecover is equivalent to running stdocimp with –nonexactclass, –noworkfloqueue, –updatesnt, and –nondeleted docs options. You must specify the –e option to request an exact document class match and the –d option to import the deleted documents from media.

Procedure

Use the following command to import both sides of a medium and place missing or damaged documents in the DOCS or DOCTABA table:

odrecover -b <surface ID>

Use the -b option if documents are missing or damaged in one table, but not the other.

To import all documents, including deleted documents, use `odrecover -d`:

odrecover -d <surface ID>

The -d flag imports the documents even if they already exist in the DOCS or DOCTABA table.

Tip

The `odrecover` program initiates a document import job. You can monitor the document recovery process like any other document import using Background Job Control.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Corporation
J74/G4
555 Bailey Avenue
San Jose, CA 95141
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims

related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names might be trademarks of IBM or other companies.

U.S. Patents Disclosure

This product incorporates technology covered by one or more of the following patents: U.S. Patent Numbers: 6,094,505; 5,768,416; 5,625,465; 5,369,508; 5,258,855.

Index

A

- advanced backup
 - cache 53
 - databases 53
 - detailed backup table 50
 - non-FileNet datasets 64
 - overview 16
 - procedures 80
 - recommended datasets 53
 - summary backup table 45
- advanced restore
 - practice backup/restore 84
 - single dataset lost 85
 - strategies 83
 - summary table 86
- AIX/6000, Image Services for
 - cache files 57
 - index database 58
 - linking databases 68
 - mksysb tape restore 65
 - Oracle control file restore 65
 - redo/transaction log 61
- archived redo or transaction log 60

B

- Background Job Control
 - consolidating media 167
 - erasing media 167
 - import log 177
- background request number 35
- backup

- daily 25, 163
- FileNet backup 84
- fnsd directory 71
- full 25
- Microsoft SQL Server databases 62
- overview 16
- practice backup 84
- raw partitions 71
- recommendations 16
- recommended datasets 25

C

- cache
 - protection 24
 - raw logical volume names 57
 - restore 162
 - restore procedures 122
 - synchronization issue 160
 - systems without optical media 163
- Capture settings 56
 - FileNet Capture settings 21
- character special files 21, 67, 72
- checkpoint file 35
- configuration database file, FileNet, after a restore 39, 87

D

- datasets
 - magnetic disk 141
 - multiple datasets lost 89
 - optical media 141

- Oracle 52
- datasets, single dataset lost 85, 89
 - cache 122, 127
 - index database 109
 - permanent database 90
 - permanent database recovery log 99
 - security database 103
 - transient database 122, 127
 - WorkFlo Queue database 114
- dbverify 99, 113
- deldocs 127, 131, 167
- deleting
 - documents 167
- disk mirroring 24
- DOC_tool 95
- document
 - IDs, preventing duplicate 180
 - importing 176
- document ID, next new 148
- duplicate documents 144, 147
- duplicate index entry 126, 130

E

- Enterprise Backup/Restore 26
- erasing
 - media 167

F

- file system 48, 50, 88
- FileNet Capture settings 56
- FileNet configuration database file 39, 87
- FileNet software 48, 50
- FileNet system tools 166
- fn_util inittrans 129

H

- HP-UX, Image Services for
 - character special files 67
 - index database 58
 - linking databases 68
 - redo/transaction log 61

I

- import
 - documents 176
- import documents 149
 - concerns 149
 - odrecover 98, 112
 - process 97, 111, 116, 121
 - stdocimp 98, 112, 117, 121
- index database
 - recovery issues 162
 - redo log, Oracle 60
 - restore 162
 - restore procedures 109
 - shutdown FileNet 109
 - transaction log, SQL Server 60

K

- killfnsw 171
- killing
 - interprocess communication
 - segments 172
 - processes 172
 - TM_daemon 172

L

- linking databases
 - examples 68
 - raw data 67

- raw logical volumes 67
- symbolic links 76
- links 21
 - character special files 67
 - examples 68
 - restore requirements for 88
 - to raw partitions 67
- logical volume groups, create 88
- logical volumes, create 88
- lost data
 - import documents 149
 - SNT_update 147
 - update scalar numbers table 146
- M**
- magnetic disk, datasets 141
- Microsoft SQL Server databases 62
- mirroring 24
- MKF
 - coordination with FileNet 155
 - coordination with third party 158
 - recovery log 154, 159
 - requirements 158
 - splicing in 154
 - recovery logs, zeroing 174
 - rollforward 154
 - trigger point 155, 159
- MKF_zeroaij 87, 101, 107, 133, 174
- mksysb tape 66, 88
- multiserver configuration
 - stopping and restarting server 136
 - Visual WorkFlo database procedure 137
- O**
- odrecover 176

- offline backup 18, 45
- operating system, lost 88
- optical media
 - datasets 141
 - importing 149
 - odrecover 98, 112
 - stdocimp 98, 112, 117, 121
 - systems without 163
- optical media surface ID 36
- update procedures 95
- used in a restore 35
- Oracle
 - control files 65
 - software 52
 - Visual WorkFlo database 137
- P**
- permanent database
 - backup 53
 - raw logical volume names 53
 - recovery log
 - backing up 53
 - initialization 99
 - restore procedures 90
 - shutdown FileNet 90
 - synchronization issue 160
 - updating scalar numbers table
 - values 146
- processes
 - killing 172
- R**
- RAID 24
- raw data partitions 20, 48
 - backup procedure 71

- names 73
- restore procedure 71
- raw disk
 - backup 27, 31
 - restore 33
- raw logical volumes, linking databases 67
- raw slices 48
- recovery log
 - MKF 158
 - permanent database 53
 - precautions 159
 - requirements 158
- redo log
 - archived 60
 - description 60
 - restore 113
- requirements
 - for a restore 88
 - for third party software 18
- rescan, batches 126, 130
- restore
 - cache 162
 - concerns 140
 - fnsd directory 71
 - index database 162
 - links 88
 - multiserver configuration 136
 - practice restore 84
 - rebuild system 88
 - requirements 88
 - synchronization problems 144
 - transient database 162
- rollforward, MKF 154

S

- scalar numbers table 146
 - checkpoint file
 - contents 146
 - editing 36
 - error condition 146
 - MKF and checkpoint values 146
 - next document ID 148
 - overview 144
 - running SNT_update 147
 - update procedures, in a restore 38, 92
- security database
 - backup 54
 - recovery log, raw logical volume
 - names 54
 - restore procedures 103
- security database recovery log,
 - initialize 105
- shutting down
 - FileNet system 171
- simple backup
 - overview 27
 - procedures 31
- simple restore
 - multiserver configuration procedures 43
 - overview 27
 - procedures 33
- snt.chkpt file 146
- SNT_update procedures 94
- SNT_update program
 - error condition 146
 - overview 147
 - procedures in a restore 38

Solaris Operating Environment, Image Services for

- cache 57
- index database 58
- linking databases 69
- redo/transaction log 61
- VERITAS 48
- special files 21, 88
- SQL Server databases 62
- surface ID 35, 36, 148
- symbolic links 67, 71
 - examples 68
 - raw data 76
- synchronization
 - cache 56, 160
 - dataset 160
 - definition 140
 - index database 162
 - magnetic disk 141
 - MKF database 154, 160
 - optical media 141
 - permanent database 160
 - problems 144
 - requirements 140
 - transient database 55, 160
- system tools, FileNet 166

T

- Task Manager 32, 34, 81, 91
- TM_daemon
 - killing 172
 - restarting 173
 - stopping in Windows Server 172
- transaction log 60, 113
- transaction log, archived 60

transient database

- backup 55
- initialize procedures 127
- raw logical volume names 55
- recovery log
 - initialize 131
 - raw logical volume names 55
- restore
 - issues 162
 - procedures 122
- shutdown FileNet 128
- synchronization issue 160

V

- VERITAS 48
- Visual WorkFlo database 137
- VW Verify 137

W

- Windows Server, Image Services for
 - cache files 57
 - default directory 51
 - index database 58
 - redo/transaction log 61
 - Registry, restore tips 66
- WorkFlo Queue database
 - backup 62
 - raw logical volume names 62
 - recovery log
 - initialize 118
 - restore procedures 114

Z

- zeroing MKF
 - recovery logs 174



Product Number: 5724-R95

Printed in USA

SC19-3327-00

