

IBM FileNet Image Services
Version 4.2

*MKF Database Migration
from Raw Partitions to File Systems
on UNIX Servers*



IBM FileNet Image Services
Version 4.2

*MKF Database Migration
from Raw Partitions to File Systems
on UNIX Servers*



Note

Before using this information and the product it supports, read the information in “Notices” on page 23.

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 2011.

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

Contents

MKF database migration to file systems on UNIX servers 7

Introduction 7

Advantages of file systems 7

Limitations 8

Document revision history 8

Accessing IBM FileNet documentation 8

IBM FileNet education 8

Feedback 8

Documentation feedback 9

Product consumability feedback 9

Strategy 9

Determining the size of the cache data sets 10

Creating the EBR Scripts 10

Output file sizes 11

Magnetic disk considerations 11

Magnetic tape considerations 12

For backing up MKF raw partitions 12

For restoring MKF file systems 14

Manually modifying the current .ddf and .res files 14

Dataset Definition file 14

Restore script 15

Automatically generating special editions of the .ddf and .res
files 16

Migrating MKF databases on multiple servers 17

Verifying the MKF databases 17

Creating links for the MKF databases 18

Backing up FileNet Capture settings	19
Backing up the MKF databases	19
Migrating from MKF raw partitions to file systems	20
Restoring the MKF databases	22
Verifying the MKF databases	22
Return to production mode	22

Notices 23

Trademarks	26
U.S. Patents Disclosure	27

Index 29

MKF database migration to file systems on UNIX servers

Introduction

IBM® FileNet® Image Services now supports either file systems or raw partitions for its Multi-Keyed File (MKF) databases on UNIX servers. File systems are the default for fresh installations. Existing FileNet Image Services users can choose to continue using raw partitions as before, or migrate to file systems.

After your current UNIX-based Image Services server has been upgraded to Image Services version 4.2, you can optionally convert the MKF databases from raw partitions to file systems. This document describes the migration process.

FileNet Image Services continues to support raw partitions for the MKF databases, if you do not choose to migrate. FileNet Image Services has always used regular files for MKF databases on Windows servers.

Advantages of file systems

File systems provide advantages over raw partitions.

- File systems do not require vendor-provided software, such as Veritas, to configure and manage the physical and logical volumes.
- File systems are less complicated to work with than raw partitions.

Limitations

All the MKF databases on a server must be either raw partitions or file systems. You cannot have a mixed combination on a single server.

If you migrate your MKF databases to file systems, you cannot convert them back to raw partitions.

Document revision history

Image Service version	Date	Comment
4.2	May 2011	Initial release.

Accessing IBM FileNet documentation

To access documentation for IBM FileNet Image Services products:

- 1 On the www.ibm.com website, enter “FileNet 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.

IBM FileNet education

IBM provides various forms of education. Please visit the IBM Information Management support page at (www.ibm.com/software/data/support).

Feedback

We value your opinion, experience, and use of our products. Please help us improve our products by providing feedback or by completing a consumability survey.

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 (<http://www-306.ibm.com/software/data/info/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.

Strategy

Plan the migration for a time when your FileNet image Services system can be off-line for about eight hours. The major steps in migrating the MKF databases from raw partitions to file systems are:

- Create FileNet Image Services Enterprise Backup and Restore (EBR) scripts to back up the databases as raw partitions.
- Create EBR scripts to restore the databases as file systems.

- Back up the MKF database raw partitions.
- Convert the MKF database definitions to file systems.
- Restore the MKF databases as file systems.

The following sections describe each of these major steps in detail.

Determining the size of the cache data sets

Before you create the EBR scripts, use the FileNet Image Services System Configuration Editor (fn_edit) to determine the size of the cache partition. You must supply this information to the EBR_genscript program later.

- 1 As the FileNet Image Services software user, start the System Configuration Editor:

fn_edit &

- 2 On the Datasets tab, locate the File Size column for cache. The partition data set size is calculated by multiplying the value in this column by 1024 (KB).
- 3 Write down the file size and calculated data set size for each cache.
- 4 Exit the Configuration Editor.

Creating the EBR Scripts

Before you perform the MKF migration from raw partitions to file systems, create the appropriate EBR scripts to back up and restore the MKF data in the correct format. You can back up to either tape or disk.

Tip Backing up the MKF databases to disk is significantly faster than backing them up to tape.

For complete information on using EBR, see the document, *Enterprise Backup and Restore*. Chapter 5, “Developing your scripts” and Appendix A, “Programs and utilities” are especially helpful. To download this document from the IBM Support web-site, see [“Accessing IBM FileNet documentation” on page 8](#).

Output file sizes

EBR backup files are compressed, and contain error correction data to increase robustness and performance.

The EBR backup file for an MKF database is typically about one half the size of the portion of the database in use; that is, the amount of data up to the high water mark.

Tip You can determine the high water mark of an MKF database by running **Xapex** > System Monitor > Storage > Databases, or by running the command-line utility **MKF_stats**.

Magnetic disk considerations

To avoid a significant degradation in performance when configuring MKF and relational databases, the data files should be on one set of disk arms, and the recovery logs should be on another set of disk arms. Otherwise, it might be impossible, after a disk crash, to restore the database and then roll the database forward using the disk-resident recovery logs.

Important Some customers current databases might be configured incorrectly—the recovery log and the data files are on the same disk arm or set of disk arms. Check for this mistake in your current configuration. If it exists, a database conversion using two sets of disk arms can correct this problem.

When using RAID, it is best to have the database data files on one RAID subsystem, and the recovery logs on another.

Magnetic tape considerations

Backing up and exporting to magnetic tape is more than an order of magnitude slower than backing up to magnetic disk files. Allot additional time for the conversion if you plan to use magnetic tape.

Tip Compute the number of tape cartridges needed. EBR backup tape cartridges must be prelabeled.

You can reduce the backup time if you have multiple tape drives. Use multiple EBR threads to back up your tape drives concurrently.

For most tape drives, EBR drives the tape at the drive's maximum throughput. Using the tape drive's hardware compression will not increase the throughput or capacity of the tape drive because the data sent to the tape drive by EBR is already compressed. Tape cartridge capacity and throughput are quoted both with and without compression turned on. Typically, the figures quoted for compression- turned- on are twice as large as the no-compression figures for cartridge capacity and tape drive throughput. When estimating your conversion times, use the figures for no-hardware-compression.

For backing up MKF raw partitions

To create an EBR script to back up your MKF databases:

- 1 Log on as a user with fnsd privileges in a shell window. You must be a user with fnsd privileges to run EBR and the platform dependent tools that manage logical volumes.

2 Start the EBR_genscript -c program:

```
Mars(user1)/fnsf/local> EBR_genscript -c

==== EBR_genscript Main Menu ====

[1] - Generate dataset definitions
[2] - Generate device specification
[3] - Generate backup script
[4] - Generate restore script

[9] - Help
[0] - Exit

Enter a command ==>3
```

3 Enter 3 at the prompt to generate a back up script.

4 Follow the prompts to create a script to back up the MKF databases on the server.

Tip In a multi-server environment, create separate scripts for the root/index and storage library servers.

5 Specify backup parameters for the current permanent, transient, and security databases.

However, the current default FileNet data set generated by EBR_genscript for the transient database cannot be reconfigured onto a file system.

6 Specify the cache backup as a partition dataset. Backing up the transient database as a normal MKF database, and backing up cache as a partition, causes the tie between them to be broken so that they can be restored later as separate file systems.

For each cache partition data set, enter the following parameters:

- Partition DatasetSize – the size you calculated earlier from `fn_edit`.
 - Partition start_block – always 0
 - Partition block_size – always 1024 bytes
 - Partition num_blocks – the DatasetSize divided by the block_size (1024) This number equals the file size you copied earlier from `fn_edit`.
- 7 When `EBR_genscript` is finished, it produces a backup script (`.bac`) that is customized to back up your MKF data set raw partitions.

Tip If you are backing up the transient database and cache the backup will have locked and ageable objects, and will add time to the conversion process.

For restoring MKF file systems

To restore the MKF databases as raw partitions, you can either make copies of your existing Dataset Definition file (`.ddf`) and existing Restore script (`.res`) and modify them, or you can use the `EBR_genscript` program to create special editions of these files.

Manually modifying the current `.ddf` and `.res` files

Dataset Definition file

Modifying the Dataset Definition file can be time consuming, especially if you need to specify one partition data set for each cache data set that is configured on the system. Your modified `.ddf` file might look, in part, like this example:

```
TranDB_cmishp02 : MKF
-- DatasetSize = 603,979,776
   location = "cmishp02";
   base_data_file = "/fnsf/dev/1/transient_db0";
end_MKF
cache0_cmishp02 : partition
   location = "cmishp02";
   filename = "/fnsf/dev/1/cache0";
-- DatasetSize = 104,857,600;
   start_block = 0;
   block_size = 1024 bytes;
   num_blocks = 102400;
end_partition
```

Notice that the `transient_db` token is not present.

Restore script

Modifying the Restore script manually requires you to add the `reconfigure_onto` option to the permanent, transient, and security databases.

The `interval_restore_follows` option must be set to false.

Your modified .res file might look, in part, like this example:

```
RESTORE_OPTIONS

PermDB_cmishp02 : restore_options
    full_restore;
    interval_restore_follows = false;
    reconfigure_onto "/fnsw/dev/1/permanent_db0";
end_restore_options

SecDB_cmishp02 : restore_options
    full_restore;
    interval_restore_follows = false;
    reconfigure_onto "/fnsw/dev/1/sec_db0";
end_restore_options

TranDB_cmishp02 : restore_options
    full_restore;
    interval_restore_follows = false;
    reconfigure_onto "/fnsw/dev/1/transient_db0";
end_restore_options

END_RESTORE_OPTIONS
```

Automatically generating special editions of the .ddf and .res files

To generate special editions of the Dataset Definition file and the Restore script:

- 1 Start the EBR_genscript program and specify the -c option:

```
Mars(user1)/fnsw/local> EBR_genscript -c

===== EBR_genscript Main Menu =====

    [1] - Generate dataset definitions
    [2] - Generate device specification
    [3] - Generate backup script
    [4] - Generate restore script

    [9] - Help
    [0] - Exit

Enter a command ==>1
```


- 2 Enter **1** at the prompt to generate a new Dataset Definition file, or enter **4** to generate a new restore script.
- 3 Follow the prompts to create the new .ddf or .res file.
- 4 Repeat the steps in this section, if necessary, to create additional files.

Migrating MKF databases on multiple servers

- 1 Stop all the servers in the proper order.
- 2 Migrate the MKF databases on the root server.
- 3 Start the FileNet software on the root server.

Leave the FileNet Image Services software running on the root server while you migrate the MKF databases on the storage library and application servers.

Verifying the MKF databases

As a safeguard, run the MKF_stats command for each of your MKF databases. Gathering the database statistics before and after the migration can help you verify that the migration was successful.

Tip To obtain the most accurate information and to avoid degrading system performance, run MKF_stats when the database is shut down.

For example, to gather the database statistics for permanent_db0, you might enter:

```
MKF_stats -d -c /fnsw/dev/1/permanent_  
db0 > PDB0_premig.txt
```

Where *PDB0_premig.txt* is the name of a text file to which the output of the MKF_stats command is sent. The file name can be anything you choose.

Repeat the MKF_stats command for each MKF database. Store the output of each command in a separate file.

Creating links for the MKF databases

The standard location for the FileNet Image Services MKF databases is the /fnsw/dev/1/ directory. To prevent performance issues, you might choose to keep some or all of the MKF databases in other locations.

The **/fnsw/etc/mkf_fs_links.sh** script creates zero-byte files and symbolic links to them. The input to this script is the fn_mkf_fs_links_input.txt file that you can customize.

Tip You can create your own text file and include it as a command-line option by following this example:
mkf_fs_links.sh custom_text_file

- 1 As the FileNet Image Services software user, locate the template file, /fnsw/etc/mkf_fs_links_input.txt file, in the /fnsw/etc/ directory.
- 2 Use your preferred text editor, such as **vi**, to modify this file with the locations of your MKF databases.

Tip List the contents of the /fnsw/dev/1 directory on each server to see the current MKF database links for comparison.

- 3 Update the file system location column of the template. For example, on an AIX combined server, you might update the template to look like the following example.

#dataset names	file system location
#-----	-----
cache0	/fs/cache
permanent_db0	/fs/perm_db
permanent_rl0	/fs/perm_rl
transient_db0	/fs/trans_db
transient_rl0	/fs/trans_rl
sec_db0	/fs/sec_db
sec_rl0	/fs/sec_rl

- At the appropriate time in the migration procedure, you are asked to run the `mkf_fs_links.sh` script. The script uses the **touch** command to create a zero-byte file in the file system location that you specified, and also adds a symbolic link statement to the `/fnsw/dev/1/` directory for each MKF database.

Backing up FileNet Capture settings

If your FileNet Image Services system uses IBM FileNet Capture, copy your configuration node from the BES to your local Capture repository before you continue. See the IBM FileNet Capture documentation for more information.

Backing up the MKF databases

- Change FileNet Image Services to backup mode.

initfnsw -y backup
- Use your EBR backup script to create an OFFLINE backup of all MKF database files and caches.

Migrating from MKF raw partitions to file systems

After you have finished the preliminary steps and backed up the MKF databases, you can schedule the migration.

To migrate the MKF raw partitions to file systems:

- 1 Start the FileNet Image Services System Configuration Editor, `fn_edit`, by entering:

`fn_edit &`

- a On the System Attributes tab, scroll to the right and toggle the **MKF Uses Raw Partitions** variable to OFF. An X must appear in the field.
- b Exit the Configuration Editor and save your changes.

- 2 Stop the FileNet Image Services software:

`initfnsw -y stop`
`killfnsw -DAy`

- 3 Rebuild the system configuration files, which include the `permanent.ddl`, `transient.ddl`, and `security.ddl` files, by entering:

`fn_build -a`

- 4 Use the **`less`** command to verify that the MKF data definition language (DDL) files in the `/fnsw/local/sd/1` directory are updated with **`file`** data type. For example, to view the `permanent.ddl` file, enter:

`less /fnsw/local/sd/1/permanent.ddl`

The contents of the permanent.ddl file look similar to the following example.

```
FILES
(
target_station "DocServer";
base data      file  "/fnsw/dev/1/fn_permanent_db0" (blocks = 12800);
recovery_log   file  "/fnsw/dev/1/fn_permanent_rl0" (start = 0, blocks =
8192);
);
```

Ensure that the data type is **file**, rather than **partition**.

- 5 If you modified the fn_mkf_fs_links_input.txt file earlier, run the fn_mkf_fs_links.sh script now to create a zero-byte file for each of the MKF databases and to create the symbolic links to those files in the /fnsw/dev/1/ directory.

/fnsw/etc/mkf_fs_links.sh

- 6 Use the fn_util command to initialize the MKF databases as file systems.

On combined servers or root/index servers:

fn_util initsec

On combined servers or storage library servers:

fn_util initperm

fn_util inittrans

Important

Do not run:

initcache – applies only to Windows servers

initnch – NCH is already stored in standard files,
not raw partitions

- 7 Change Image Services to restore mode

initfnsw -y restore

Restoring the MKF databases

Use the EBR Restore script that you created earlier to restore the MKF database files from your EBR backup. The MKF databases are restored as file systems.

Verifying the MKF databases

Earlier in this procedure, you ran the MKF_stats command to gather the statistics for each MKF database. Now that the migration is finished, run the MKF_stats command again so you can compare the database statistics before and after the migration and verify that the migration was successful.

For example, to gather the database statistics for permanent_db0, you might enter:

```
MKF_stats -d -c /fnsw/dev/1/permanent_  
db0 > PDB0_postmig.txt
```

Where *PDB0_postmig.txt* is any name you choose for a text file to which the output of the MKF_stats command is sent.

Repeat the MKF_stats command for each MKF database, and store the output of each command in a separate file.

Compare the *premig.txt* and *postmig.txt* files for each MKF database to verify that the migration was successful.

Return to production mode

The MKF database migration to file systems is complete and you can return your FileNet Image Services system to production mode.

```
initfnsw -y stop  
killfnsw -ADy  
initfnsw start
```

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 Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Microsoft and Windows 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 may be trademarks or service marks of others.

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

C

commands

 fn_build 20

 fn_edit 20

 less 20

 MKF_stats 17, 22

configuration editor 20

Converting the MKF database

 on multiple servers 17

E

EBR 12

Exporting

 to tape 12

F

File sizes

 output 11

fn_build command 20

fn_edit command 20

L

less command 20

M

Magnetic tape

 considerations 12

MKF databases

 verifying 17, 22

MKF_stats command 17, 22

O

Output file sizes 11

R

RAID 11

S

system configuration editor 20

V

verifying MKF databases 17, 22



Product Number: 5724-R95

Printed in USA

SC19-3311-00

