IMAGE SERVICES

System Configuration Editor Overview

**IS 4.0 HP Integrity Edition
and IS 4.0 SP5**

9844079-004

November 2006

Contents

2   Online Help Text   28

About This Manual

This *System Configuration Overview* describes the FileNet tools for viewing and editing your Image Services configuration. It is designed primarily for field personnel to use when setting up customer Image Services configurations.

For a printed version of the online help text, see **Chapter 2, "Online Help Text," on page 28**.

Conventions Used in this Manual

The following paragraphs discuss the ways in which we call your attention to information throughout this document.

File Paths

Since this manual is used for all platforms, examples of file path designations, where used, are given for both UNIX and Windows platforms. For example:

/fnsw/local/sd   (UNIX)

\fnsw_loc\sd   (Windows)

Console Displays

Information you see displayed at your console is shown in this document in the following manner:

```
Surface '3176' : 1 document processed
Local doc_id = '2235007' Original doc_id = '2235007'
Original ssn = '8502'
Primary copy. No tranlog copy exists.
* document successfully deleted from databases. *
* Purging pages from disk... *
* This document has been successfully purged. *
```

Cautions, Notes, and Tips

Three message types call your attention to important information:

CAUTION Signals possible damaging consequences of an action, such as loss of data or time.

Note Draws your attention to essential information you should be sure to read.

Tip Introduces an idea that might make your work easier.

Command Syntax

Command syntax definitions are indented and in bold text:

ddexim –e > <filename>

Optional Parameters

Optional parameters and keywords are within square brackets:

ddexim [–e] > <filename>

Required Parameters

Parameters that require you to provide information are shown within angle brackets (< >).

For example, for the following command:

ddexim –e > <filename>

you must substitute the name of a command for the parameter in angle brackets, such as:

ddexim –e > myfile

FileNet Education

FileNet provides various forms of instruction. Please visit Global Learning Services on FileNet's Web site at: **www.filenet.com**.

Comments and Suggestions

FileNet invites all customers to communicate with the Documentation group on any question or comment related to FileNet manuals and online help. Send email to **docs@filenet.com**. We will make every effort to respond within one week. Your suggestions help us improve the products we deliver.

1

System Configuration Tools

System Configuration Tools are the utilities you use to configure the FileNet Image Services software for your system. System Configuration Tools work in conjunction with the FileNet Task Manager. Task Manager provides a graphical user interface (GUI) and online help.

System Configuration Tool Programs

The System Configuration includes three programs:

fn_edit (FileNet Image Services - System Configuration Editor)
fn_build (FileNet Configuration Builder)
fn_util (FileNet Configuration Utilities)

fn_edit

The **fn_edit** program invokes a graphical user interface (GUI) utility called the *FileNet Image Services - System Configuration Editor*, which enables you to create, display, and modify the system configuration database.

With the FileNet Image Services - System Configuration Editor, you can easily configure a new system using a series of dialog boxes. When you create a new Configuration Database, the initialization procedure guides you through the configuration process appropriate for the server environment you specify. When you finish the procedure, the remaining configuration parameters are set to suitable default values. You must manually change any configuration values that differ from the defaults.

The System Configuration Editor also has online help. The Online help provides a table of contents and keyword searches, hypertext links, and browse sequences.

The Help pull-down menu contains these options:

Help Menu Option	Function
Contents	Display table of contents for online help
Search for help on	Search for a particular help topic
How to use help	Instructions on using online help
On Context Help	Display a help window related to the currently-active window
About	Display information about the current version of Image Services software

For more information on the System Configuration Editor, see:

- **"FileNet Image Services - System Configuration Editor" on page 16**

- **"How to Use the FileNet Image Services – System Configuration Editor" on page 23**

- **"System Configuration Editor Online Help" on page 28**

fn_build

The **fn_build** program creates various configuration files. The fn_build program is run automatically by the FileNet Task Manager when the FileNet software is restarted. You can also run fn_build from a command prompt.

For more information on fn_build, see **"Configuration Databases" on page 13**.

fn_util

The **fn_util** program starts, stops, creates, and updates the MKF and RDBMS databases, calling fn_oracle, fn_db2, and fn_mssql to perform RDBMS configurations. You run fn_util from a system prompt on UNIX servers, or a command prompt on Windows.

For more information on fn_util, see .

Configuration Databases

The fn_build, fn_edit, and fn_util programs use the Configuration Database (CDB) abstract. Applications obtain configuration information by making calls to CDB. In cases where getting configuration information from CDB is not practical, applications can read configuration files from the directory /fnsw/local/sd (UNIX) or \fnsw_loc\sd (Windows).

The Image Services Configuration Database is created during system initialization or migrated to a new revision level during an IS software upgrade. It resides in /fnsw/local/sd/conf_db (UNIX) or \fnsw_loc\sd\conf_db (Windows), and is named:

IMS_xxx.cdb

where xxx is the version number and cdb stands for Configuration Database. Each time you run the fn_edit program, and save it, the version number is incremented. (The current version of the cdb is saved, and a new version is created. This new version has the version # incremented in the name.) To support multiple server configurations, the configuration database can be remotely accessed.

The system configuration file is in ASCII format, so utilities to import and export the file are not necessary. Other configuration files like

as_confg and print_config are generated by fn_build from information in the Configuration Databse.

System Configuration also includes eight read-only databases that reside in /fnsw/lib/conf_db (UNIX) or \fnsw\lib\conf_db (Windows). These five files are described in the following table.

Configuration File	Description
FN_CMB_1.CDB	Read-only database template for Combined servers. Contains the default values used to create a new database.
FN_DUO_1.CDB	Read-only database template for Dual Server con-figuration.
FN_RES_1.CDB	Read-only database template for Remote Entry Server configuration.
FN_VWF_1.CDB	Read-only database template for Visual WorkFlo Server configuration.
FN_DEV_1.CDB	Read-only database containing information about devices.
FN_EP_1.CDB	Read-only database containing information about edit procedures.
FN_PRT_1.CDB	Read-only database containing information about printers.
FN_MIG_1.CDB	Read-only database containing information about migration.

The schema of a CDB is embedded in the data file; when the schema changes, the template data file must be updated with new schema data. When the schema is updated, the revision level is incremented. For example, the CDB schema is at revision level 23 in IS Release 3.6 ESE, level 34 in IS Release 4.0, and level 43 in IS Release 4.0 SP3.

The update process uses fn_migrate to determine the revision level of a CDB file and upgrades it to the current level, supplying default values for new parameters and deleting parameters that are no longer needed. The migration specification file, FN_MIG_1.cdb, is installed in the directory /fnsw/lib/conf_db (UNIX) or \fnsw\lib\conf_db (Windows).

The update process also places an fnsw_update trigger file in the following directory:

/fnsw/local/tmp/fnsw_update (UNIX)

\fnsw_loc\tmp\fnsw_update (Windows)

So when the **initfnsw start** command is issued to start Image Services after a software update, and the trigger file exists, theTM_ daemon calls fn_migrate to update the CDB schema.

The schema update occurs only once for each format change, takes only a few seconds, and is transparent to other processes and to users.

Administration Tools Restricted to fnadmin

To execute FileNet programs that modify configuration information, modify databases, or start or stop the Image Services software, you must be logged in as a member of the fnadmin group. The root logon is a member of fnadmin.

Use of this group allows a system administrator to differentiate between users of end-user server applications, such as those available from the Applications menu, and FileNet administrative users.

The release tools create the fnadmin group.

The following programs validate, at startup, that the user is a member of the fnadmin group.

fn_build
fn_dataset_config
fn_edit
lic_admin
nch_tool
nch_update
task_man

FileNet Image Services - System Configuration Editor

The FileNet Image Services - System Configuration Editor program (fn_edit) is started from a system prompt on UNIX servers, or from the Start Menu on a Windows server. Before edits can take effect, you must use the Task Manager to restart the FileNet software, which automatically runs fn_build.

FileNet Image Services - System Configuration Editor Overview

The central element of the system configuration tool set is the CDB abstract. (CDB stands for Configuration Database). The GUI editor, fn_edit (System Configuration Editor), creates and modifies the system configuration database, IMS_xxx.CDB, via calls to CDB. The other programs, fn_util, fn_build, and Task Manager, access configuration data via CDB to control their actions.

The MKF and relational databases are started, stopped, created, and updated by fn_util. For relational database commands, fn_util calls fn_oracle, fn_db2, or fn_mssql to carry out the command.

The configuration files controlling some of the programs started by Task Manager's tm_daemon (the GUI front end to the FileNet Image Services software) are created by fn_build.

Task Manager determines which programs to run based on the SERVER_Processes table accessed from CDB. Notice that Task Manager runs fn_build. When a configuration change is made, you only need to restart the software, and the configuration files are generated automatically.

FN_DEV_1.CDB, FN_EP_1.CDB, and FN_PRT_1.CDB files are distributed read-only databases that contain information about devices, edit procedures, and printers respectively. FN_CMB_1.CDB is a database template for combined servers and contains the default values used when a new database is created. Other template files include:

FN_DUO_1.CDB (Dual Server configuration),
FN_RES_1.CDB (Remote Entry Server configuration), and
FN_VWF_1.CDB (Visual Workflo Server configuration).

These files reside in the /fnsw/lib/conf_db directory.

Task Manager and System Configutation Editor are documented with online help. Both fn_util and fn_build output help when called without arguments and are started from a system prompt on UNIX servers or command prompt on Windows. The *Software Installation Procedure* documents contain detailed instructions for running each of these programs when a new system is created.

Procedures tab

The System Configuration Editor window is designed in a file folder format with each tab of the folder containing a major subset of editor functions.

Unless the system configuration database is opened in read-only mode, the procedures tab displays first as shown on the next page.

To run a procedure, select one from the list and click the *Run* button. Dialog boxes will then appear and prompt you for the information you must provide to complete the task. You can select choices from list boxes, type in new information, or accept the default entries.

The following procedures are listed on the Procedures tab:

Add a printer
Delete a printer
Automatically configure optical storage library
Manually configure optical storage library
Delete an optical storage library
Configure an MSAR storage library
Delete an MSAR storage library
Add an optical storage drive
Delete an optical storage drive
Add a tape drive
Delete a tape drive
Set to site-control *
Set to FileNet-control **

 * Only appears if Oracle is FileNet-controlled
** Only appears if Oracle is site-controlled
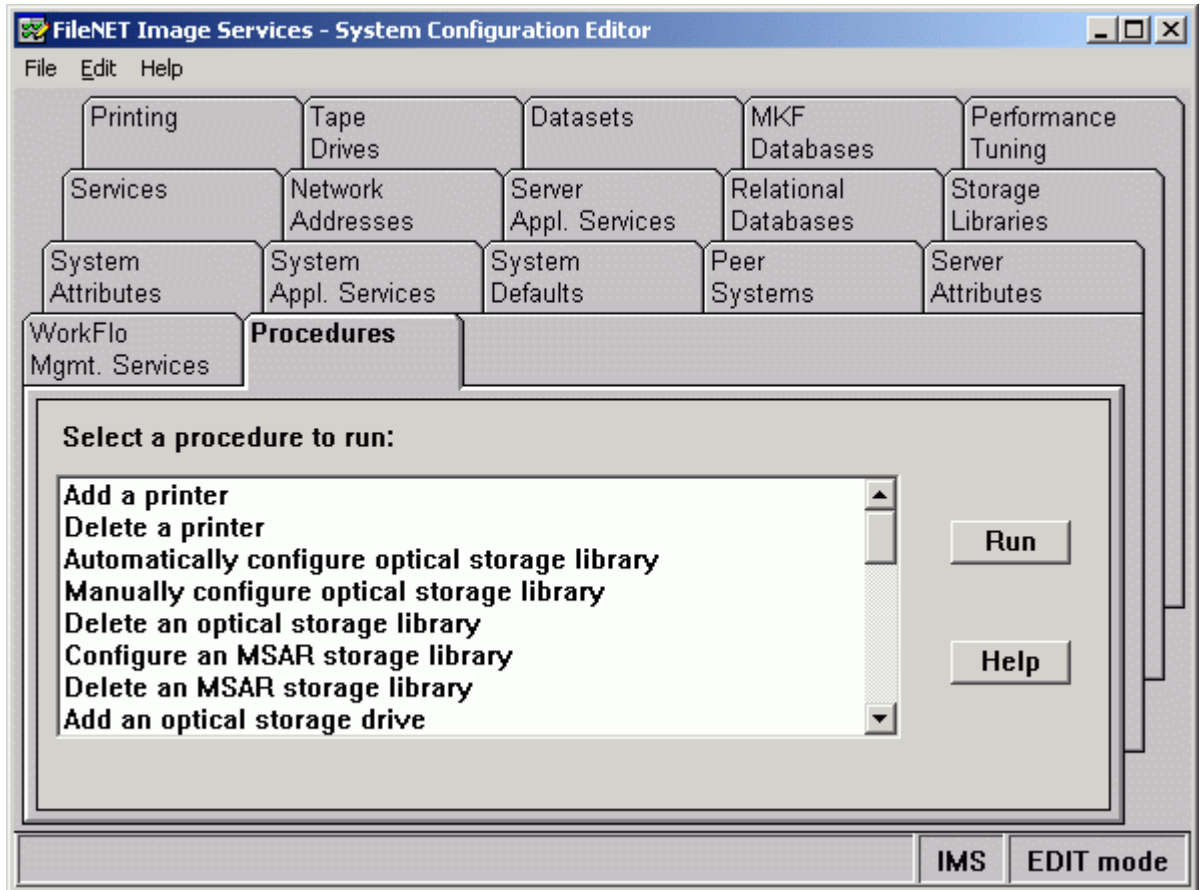
(continued from previous page)

Add Relational Database Object
Delete Relational Database Object

(continued from previous page)

Set up required Oracle datasets on root/index server *
Deconfigure Oracle datasets on root/index server *
Set up required M/S SQL datasets on root/index server **
Deconfigure M/S SQL datasets on the root/index server **
Create an optional relational DB dataset
Delete an optional relational DB dataset
Add an additional dataset
Delete an additional dataset
Add a peer system
Remove a peer system
Add an optional logical cache
Delete an optional logical cache
Add an Application Server
Delete Server
Add a Service to a Server
Delete a Service from a Server
Add a Storage Library Server
Add VWService
Modify VWService
Delete VWService

 * Only appears if you have an Oracle RDBMS.
 ** Only appears if you have an Microsoft SQL Server RDBMS

Other System Configuration Editor tabs

Besides the Procedures tab, there are several other tabs where you can view, and possibly modify, configuration data.

Each System Configuration Editor tab you select displays a different window. For example, if you click the Datasets tab the following screen displays.

Server Name	File Name	File Size (MB)
sunlight	/fnsw/dev/1/oracle_tr0	400
sunlight	/fnsw/dev/1/oracle_sys0	200
sunlight	/fnsw/dev/1/oracle_rl1	20
sunlight	/fnsw/dev/1/oracle_rl0	20
sunlight	/fnsw/dev/1/oracle_db0	200
sunlight	/fnsw/dev/1/cache0	100
sunlight	/fnsw/dev/1/permanent_db0	100
sunlight	/fnsw/dev/1/permanent_rl0	64
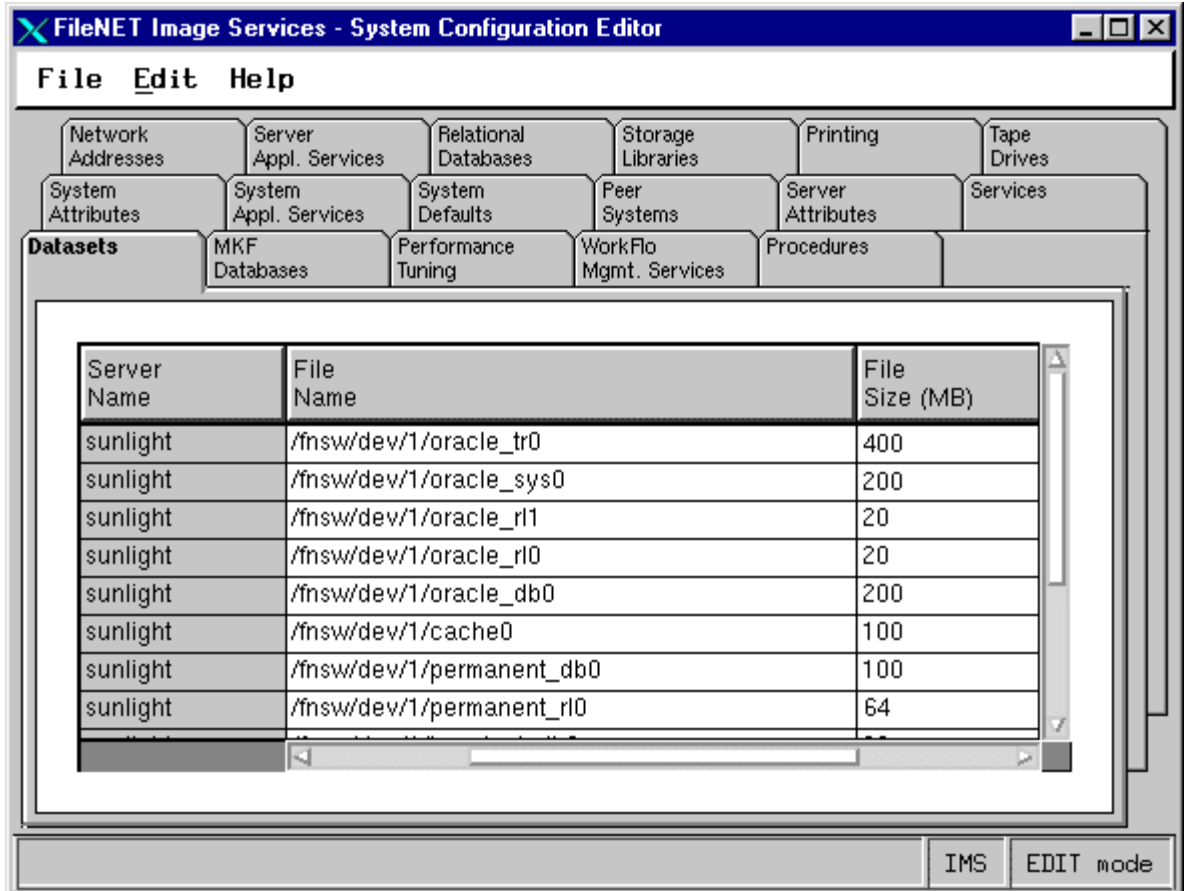
The system in this example has an Oracle RDBMS.

Tailored Interfaces

The System Configuration Editor tabs, procedures, and columns are displayed, hidden, or read-only, based on a table keyed to the product type in the software license. For example, only WorkGroup configuration parameters display on a WorkGroup Image Services system.

The System Configuration Editor includes additional changes to tailor the interface for a WorkFlo server:

- You cannot access Database Maintenance, Storage Library Control, Background Job Control, or Cache Backup from the Applications menu.

- In System Monitor, you cannot select the Magnetic Disk Cache and Storage Libraries options on the Storage menu or the Document Services Statistics option on the Statistics menu.

- You see only the appropriate Configuration Editor parameters; some tabs and columns do not display and you cannot run certain procedures.

- The SCouT Capacity Planning Tool provides only the appropriate WorkFlo server entries.

Customized Startup and Shutdown Scripts

The TM_daemon interprets the SERVER_Processes table generated by the System Configuration Editor (fn_edit). You cannot permanently edit the SERVER_Processes table since the table regenerates itself whenever you create a new IMS_xxx.cdb file by making and saving a change to an fn_edit tab.

To allow you to run customized scripts or batch files at startup or shutdown, TM_daemon looks for customized startup and shutdown files on the local drive. If a file is found, fn_edit inserts the startup file at the end

of the Image Services startup mode and inserts the shutdown file at the beginning of the Image Services shutdown mode.

Locate customized files for Image Services as follows:

• UNIX servers:

> **/fnsw/local/sd/ims_start**
> **/fnsw/local/sd/ims_stop**

• Windows servers:

> **\fnsw_loc\sd\ims_start.bat**
> **\fnsw_loc\sd\ims_stop.bat**

Customized scripts are run in the foreground unless the script itself places an operation in the background.

How to Use the FileNet Image Services – System Configuration Editor

The FileNet Image Services – System Configuration Editor consists of a series of editing tabs where you can view and edit your system configuration. In addition, File and Edit menus are located at the top of the System Configuration Editor window.

Changes to your system configuration are made by clicking the tab that contains the configuration data, and then the cell within the tab that has the value you wish to change. Depending on the data you wish to change, you can either select from a pulldown list or type in a new value.

Note   Some tabs allow you to change the state of an item by marking it with a check mark or with an X. A check mark indicates a Yes (available or

enabled) response, and an X indicates a No (unavailable or disabled) response.

The following is a list of the System Configuration Editor editing tabs. If you want to find Help information on a specifc tab you can go directly to the page number listed, or if you are viewing this document Online simply click on the tab you want below.

fn_util Program

The **fn_util** program starts, stops, creates, and updates the MKF and RDBMS databases, calling fn_oracle, fn_db2, and fn_mssql to perform RDBMS configurations. Run fn_util from a System Prompt on UNIX servers, or a Command Prompt on Windows Server.

The syntax is:

fn_util [command]

Running fn_util with no arguments displays the fn_util help text.

The commands and their actions are included in the following table.

fn_util Command	Action
abortrdb	shutdown relational database
check_osconfig	check for minimum O/S requirements to run fnsw (AIX, HP-UX only)
cinitrdb	initialize relational database for RDBMS coexistence
configrdb	configure relational database (Windows only)
init	initialize all datasets below (not for conversions)
initsec	initialize security database
initperm	initialize permanent database
inittrans	initialize transient database and cache
initrdb	initialize relational database
initnch	initialize nch database
initconv	initialize empty relational database for conversions
linkrdb	creates links to Oracle shared library (UNIX only)
mk_links	create device and dataset symbolic links (UNIX only)

fn_util Command	Action, continued
remove_datasets	removes datasets and registry entries made by fn_oracle
servicerdb	update IMS Service dependency for RDB services (Windows only)
spacerdb	initialize relational udb and tr databases
start	starts
startdb	start all databases
startsec	start security database
startperm	start permanent database
starttrans	start transient database
startrdb	start relational database
stopdb	shutdown all databases
stopsec	shutdown security database
stopperm	shutdown permanent database
stoptrans	shutdown transient database
stoprdb	shutdown relational database
update	update all datasets below
updatesec	update security database
updateperm	update permanent database
updatetrans	update transient database
updatench	update nch database
updatecache	update cache partitions (Windows only)
updaterdb	update relational database
whichfn	print release information

Re-initializing the Permanent Database

If you ever have to re-initialize your permanent database, using **fn_util init** or **fn_util initperm**, you may see the following error message:

```
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.
```

If this occurs, you must run SNT_update to resolve the problem. For information on running SNT_update, refer to the *Image Services Third-Party Backup/Restore Guidelines*.

2
Online Help Text

This chapter provides an overview of the System Configuration Editor Online Help text. The text has been modified slightly to clarify the relationship of context-sensitive help text segments. Future Image Services releases may not include printed versions of the online help text.

System Configuration Editor Online Help

The following is a list of the System Configuration Editor editing tabs. To view the Online Help for a specific tab you can go directly to the page number listed, or if you are viewing this document Online simply click on the tab you want below.

- **"Services Tab" on page 74**
- **"System Appl. Services Tab" on page 75**
- **"System Defaults Tab" on page 81**
- **"System Attributes Tab" on page 83**
- **"Tape Drives Tab" on page 85**
- **"WorkFlo Mgmt. Services Tab" on page 86**

Datasets Tab

This tab enables you to view and modify characteristics of the datasets currently configured on this server. To delete datasets or to create additional or optional datasets, use the Procedures tab.

Server Name

This field displays the host Name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Dataset Type

This field displays the type of the dataset, such as Perm_DB, Trans_DB, Sec_DB, etc. This field is locked and cannot be edited. To initialize all of the dataset names or to add/delete a specific dataset, go to the Procedures tab.

File Name

The full path name for the dataset is shown in this field. This is the device driver pathname for the dataset (for example, d:\fnsw\dev\1\cache0). To create additional or optional datasets, go to the Procedures tab.

File Size (MB)

This field contains the current size of the dataset in megabytes.

Note The minimum sizes on your server could be slightly different, depending on the server type and its configuration.

The following minimum dataset sizes are required:

Dataset	Minimum Size
Cache	100 MB
Transient Database	40 MB
Transient Recovery logs	64 MB
Permanent Database	100 MB
Permanent Recovery logs	64 MB
* Oracle System Tablespace	200 MB
* Oracle Database	200 MB
* Oracle Recovery logs	20 MB
* Oracle User DB	200 MB
* Oracle Temp Space	400 MB
MS SQL Server Temp Space	400 MB
MS SQL Server Redo Log	20 MB
MS SQL Server Database	200 MB
MS SQL Server User Database	200 MB
Security Database	64 MB
Security Redologs	64 MB

* Displayed for FileNet-controlled Oracle only

Note The Oracle User DB is optional and will only appear on the Datasets tab if one has been created.

The number displayed in the previous table is the minimum size of the dataset partition in megabytes. You can increase the size of each dataset type, if necessary. Cache is the only dataset that you can reduce in size (after clearing all caches).

If you choose to change these values, use the following increments for modifying each dataset type:

Dataset	Increment
Cache	100 MB
Transient Database	40 MB
Transient Recovery logs	64 MB
Permanent Database	100 MB
Permanent Recovery logs	64 MB
* Oracle System Tablespace	200 MB
* Oracle Database	200 MB
* Oracle Recovery logs	40 MB
* Oracle User DB	200 MB
* Oracle Temp Space	400 MB
MS SQL Server Temp Space	400 MB
MS SQL Server Redo Log	20 MB
MS SQL Server Database	200 MB
MS SQL Server User Database	200 MB
Security Database	64 MB
Security Redologs	64 MB

* Displayed for FileNet-controlled Oracle only

Note The increment sizes on your server could be slightly different, depending on the server type and its configuration.

MKF Databases Tab

The MKF Databases tab contains the following fields:

Server Name

This field displays the host Name for the server. The name is created initially by the setup program when the software is installed. This field is locked and cannot be edited here. To change the server name, go to the Server Attributes tab.

Database Type

This field is locked and cannot be edited. It displays the database type and can be either Permanent, Transient, NCH, or Security.

Max Concurrent Transactions

This number determines the maximum number of transactions that will ever be concurrently executed by MKF. The number of transactions allowed depends on two factors: the size of the database and the size of the data pool.

Max Concurrent Long Transactions

This number determines how many of the Maximum Concurrent Transactions may be long transactions. The value must between the default value and the number of Max Concurrent Transactions.

Verify Disk Writes

This switch provides protection against disk I/O problems by reading back each block and verifying it after writing to it. This option is not useful if the database is not in a raw disk partition. Click in this field to toggle between Yes and No.

Overwrite Recovery Log Action

This field contains three choices: Warning Message, Stop, or Ignore. Click the down arrow and select the action you want.

Block size

This field contains the block size (KB) of the specified MKF database. The default value is 8 KB for the Permanent, Transient, and Security MKF databases.

Recovery Log Buffer Blocks

This field contains the number of database blocks (of the above specified size), that are allotted to the recovery log buffer. The default value is 24.

Network Addresses Tab

The Network Addresses tab contains the following fields:

Server Name

This field displays the host Name for the server. The name is created initially by the setup program when the software is installed. This field is locked and cannot be edited here. To change the server name, go to the Server Attributes tab.

Network Name

This field can only be used by customers running client Image Services Toolkit (ISTK) software version 3.6 or higher. Earlier versions of client ISTK software will ignore this field and instead retrieve a Network Address.

This field allows a newer client software to resolve a standard DNS name into an appropriate Network Address for the client. This can be very useful for clients that are on different network

segments than the server.

This field is also optional. You do not have to specify a Network Name in this field if you enter a Network Address. The newer client WAL software will attempt to retrieve a Network Name. If it fails, it will retrieve a Network Address instead.

To support mixed versions of client WAL software, you can enter both a Network Name and a Network Address. New clients will request the Network Name, older clients will request the Network Address.

The Network Name you enter in this field is the DNS host name of the server. The name must be unique, less than 256 characters, and composed of alpha, digits, dot, dash, underline characters only (No Spaces).

Network Address

There are eight Network Address fields that are available for different TCP/IP network protocol (IP) addresses. Each field displays the four-part TCP address for the corresponding server name. You can add or modify network addresses in any of these fields.

The TCP address must be a four part number. Each part must be between 0 and 255 and separated by a period. (for example, 135.0.20.91).

Storage Libraries Tab

The Storage Libraries tab enables you to view and modify parameters for your system's storage libraries and storage drives. Parameters for each are found on their respective subtabs under the Storage Libraries tab.

Storage Libraries Subtab

Storage Libraries is a subtab of the Storage Libraries tab. This subtab contains the following fields:

Server Name

This field displays the host Name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Library Name

Values in this field are automatically assigned by the system and correlate to the library number (for example, Library number 0=Library A, Library number 1=Library B). Library numbers are assigned when configuring a storage library. Go to the Procedures tab to configure a storage library.

Library Type

Use the pulldown menu arrow to choose the type of storage library connected to your system.

Note For MSAR, the model number refers to the number of virtual slots in the library. i.e., Model 128 has 128 slots, Model 256 has 256 slots, etc. Furthermore, if you change the library type from a lower model number to a higher model number, for example from Model 128 to Model 256, and save this change, you will not be able to change back to a lower

model number. In fact, model numbers lower than what is currently listed in the Library Type list box, will not even be available in the drop-down list.

Device Address

This is the SCSI hardware address for your storage library. It is usually a four part number, each part separated by a space. The first number designates the bus ID, the second designates the port ID, the third is the storage library device ID, and the fourth is the logical unit number. If a library is connected to an IBM 5712 or 5736 adapter card, the library's SCSI address is a five part number, with the first two numbers designating the bus ID.

To determine the SCSI parameters for the storage libraries connected to your system, enter:

ls /dev/fnsod* (on UNIX servers)

fndev -v (on Windows servers)

Note Since MSAR libraries aren't actual hardware devices, they do not require or use a device address. Therefore, this field is disabled for MSAR storage libraries. The number will always be 0 0 0 0.

MSAR Surface Creation Directory

This column contains the directory location where all created MSAR surface files will be placed for each MSAR library. If you need to, you can edit this field to change the location path, or name, of the directory. The path can be up to 241 characters in length.

Note For optical storage libraries, this field is not enabled.

Storage Drives Subtab

Storage Drives is a subtab of the Storage Libraries tab. This subtab contains the following fields:

Server Name

This field displays the host Name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Library Name

Values in this field are automatically assigned by the system and correlate to the library number (for example, Library number 0 = Library A, Library number 1 = Library B). Library numbers are assigned when configuring a storage library. Go to the Procedures tab to configure a storage library.

Library Type

This designates the type of storage library connected to your system. Go to the Storage Libraries tab to modify this field.

Drive Number

This designates an ID number for the storage drive. Acceptable values are 0 through 12. The ID numbers are normally assigned when a storage library is configured on the system.

Storage Drive Type

This is a the type of storage drive used in the library. It can be an optical or MSAR drive. To identify an optical storage drive, you need to know the manufacturer, the optical disk size, the capacity and/or its model number. For MSAR drives, there is only one choice.

Device Address

The SCSI address of the storage drive is a four part number composed of the following elements:

- bus ID

- port ID

- device ID

- logical unit number

For a drive that is connected to an IBM 5712 or 5736 adapter card, the SCSI address is a five part number.

- bus ID (part 1)

- bus ID (part 2)

- port ID

- device ID

- logical unit number

Note   Since MSAR libraries aren't actual hardware devices, they do not require or use a device address. Therefore, this field is disabled for MSAR storage libraries. The number will always be 0 0 0 0.

Favor Writes

The default value for this field is No. Normally, read requests are given priority over write requests. However, if you want this drive to give priority to write requests for fast batch committal, click on the field.

Peer Systems Tab

The Peer Systems tab contains the following fields:

Domain Name

This field displays the Network Clearinghouse domain name for a peer system. Peer systems are used to represent both the remote and compatible system concepts.

A remote system is recognized by its internet address, is accessable via a network, and is usually a recipient for storage media data from a remote entry station (which has no storage library).

A compatible system is set up with a unique range of document and surface IDs. This allows an easy import of documents and media from one system to the compatible system. Compatible systems are usually not on the network and are recognized by their system serial number.

You must identify the internet address for each remote system and the system serial number for each compatible system when designating them as a peer.

System Serial Number

This is the serial number of the remote or compatible system (mandatory for a compatible system).

Network Address

This is the four part internet address (for example, 135.20.20.52) of the remote or compatible system (mandatory for a remote system).

Performance Tuning Tab

Performance tuning options allow you to adjust a variety of system parameters to suit the use and configuration of your FileNet Image Services system. These parameters are located on three subtabs.

System Processes Subtab

System Processes is a subtab of the Performance Tuning tab. This subtab contains the following fields:

Domain Name

This field contains the name of your system's domain. The name shown here is used as the "domain" part of your system's NCH name.

DS Notify Processes

This field contains the number of Document Services notify processes. The default value is 2.

RMT Commit Processes

Number of RMT commit processes. This value must be greater than 0; the default value is 1.

Number of dtp tran processes

The dtp tran process is responsible for cataloging documents in the Index database using Document Service's fast batch committal option. Under high volume committal scenarios, the dtp tran process can become bottlenecked.

Performance can be improved, therefore, if more than one dtp_tran process is running on a system.

The default value for this field is one dtp_process. You may select a value of 1 - 16.

CAUTION Increasing the number of dtp tran processes should only be done on systems which have multiple CPU's. In addition, the number of dtp processes should never exceed the total number of CPU's on your system.

Server Processes Subtab

Server Processes is a subtab of the Performance Tuning tab. This subtab contains the following fields:

Server Name

The name of the server whose configuration database you are currently editing.

CSM I/O Process Count

Cache Manager Process Count is the number of I/O worker processes doing asynchronous magnetic disk cache I/O. The maximum usable value is the minimum of the number of storage media and magnetic drives on the system. The default value is 2.

DTP Process Count

Data Transfer Process Count is the number of DTP processes per drive. You can select a value of 1 or 2. The default is 1.

Note The DTP Process Count value affects OSAR drives only. For MSAR, the DTP Processs Count is always 1 regardless of the number in this field.

BES Commit Process Count

BES Commit Process count is the number of bes_commit processes started when the server boots. You can start a maximum number of four bes_commit processes on a server. The default value is 2.

Server Memory Subtab

Server Memory is a subtab of the Performance Tuning tab. This subtab contains the following fields:

Server Name

This field contains the name of the server whose configuration database you are currently editing.

Document Buffer Count

This field contains the number of document and descriptor buffers to be allocated by the DBP abstract. The number must be between 1 and 128. The default value is 16. A smaller number uses less memory, but a larger number provides better performance.

Note For MSAR Storage Library systems, use the following configuration guidelines to calculate the Document Buffer Count for each applicable server: 1 per MSAR drive, +1 per optical drive, +1 per scanner, +1 per COLD process, +1 per HP11 or any other document entry program, +16.

Document Buffer Size (KB)

This field contains the size of the document buffer(s). The number must be between 16 and 1024 KB. The default is 64 KB.

More complex images require a larger buffer size, while less complex images do not. You may be able to reduce the buffer size and increase the number of buffers to increase system performance; or you may want to increase the buffer size to accommodate extremely complex images. If necessary, use the **dbp –s** command to display buffer size and count statistics for fine tuning the system.

Directory Buffer Count

This field contains the number of directory buffers to be allocated by the DBP abstract. The number must be between 1 and 64. The default value is 4. A smaller number uses less memory, but a larger number provides better performance.

Directory Buffer Size (KB)

This field contains the size of the directory buffer(s). The number must be between 1 and 256. The default is 16 KB.

If memory is limited, you can change the buffer size to two times the number of fast batch committal jobs that will be running concurrently. For example, if you specified (on the Performance Tuning/Server Processes tab) that four fast batch committal jobs (bes_commit processes) would run on the server, the buffer size you enter here would be 8 KB (2 x 4).

Oracle Cache Size (MB)

(Filenet-controlled Oracle only)
This field contains the size of the Oracle database cache. The default size is 100MB.

A larger number could improve performance, depending on how much memory is available and many other factors. Since Oracle performance tuning issues are complex, refer to the

Oracle documentation for more information on setting this parameter.

Oracle Log Buffer Size (bytes)

(Filenet-controlled Oracle only)

This field contains the size of the Oracle redo log buffer. This is the portion of the Oracle redo log cached in memory. The default size is 163840 bytes.

A large number could improve performance, depending on how much memory is available and many other factors. Since Oracle performance tuning issues are complex, refer to the Oracle documentation for more information on setting this parameter.

Permanent Buffer Pool Size (KB)

This field contains the size of the permanent database buffer pool. This is the portion of the permanent database cached in memory. The default size is 2500 KB. To increase performance by having a small database entirely resident in memory, the size of this buffer should be almost twice the size (in KB) of the data and index (but not the description) in the database.

Note After making any change to this field, additional steps **must** be performed. See the section, **"Increasing MKF Buffer Pools / Shared Memory" on page 46**, for more information.

Transient Buffer Pool Size (KB)

This field contains the size of the transient database buffer pool. This is the portion of the transient database cached in memory. The default size is 2500 KB. To increase performance by having a small database entirely resident in memory, the

size of this buffer should be almost twice the size (in KB) of the data and index (but not the description) in the database.

| **Note** | After making any change to this field, additional steps **must** be performed. See the section, **"Increasing MKF Buffer Pools / Shared Memory" on page 46**, for more information. |

Security Buffer Pool Size (KB)

This field contains the size of the security database buffer pool. This is the portion of the security database cached in memory. The default size is 2000 KB. To increase performance by having a small database entirely resident in memory, the size of the buffer should be almost twice the size (in KB) of the data and index (but not the description) in the database.

| **Note** | After making any change to this field, additional steps **must** be performed. See the section, **"Increasing MKF Buffer Pools / Shared Memory" on page 46**, for more information. |

Max Shared Memory (MB)

This field contains the maximum amount of server shared memory allowed with the current configuration. The default memory sizes are 160 MB for AIX platforms, and 2048 MB for all others.

When Image Services starts, it only allocates shared memory as needed, so the actual amount of shared memory in use could be far below the value in this field. If the number in this field is not enough to handle the needs of the customer, then this value should be increased as needed.

Note After making any change to this field, additional steps **must** be performed. See the section, **"Increasing MKF Buffer Pools / Shared Memory" on page 46**, for more information.

Increasing MKF Buffer Pools / Shared Memory

CAUTION Changing the MKF Buffer Pool sizes or Shared Memory size impacts system performance. Modifying these fields should only be done with the advice and consent of FileNet CSS personnel. Be sure to make small adjustments to these field sizes over a period of time so their effect on the system can be monitored.

Follow these steps if you have increased the Buffer Pool Size (KB) field(s) for one or more of the datasets and/or the Max Shared Memory (MB) field.

1 Exit **fn_edit** and save your changes.

2 Enter the following command to rebuild the system configuration files:

fn_build -a

Make sure the command runs successfully by checking that no errors have occurred.

3 If you changed only the Max Shared Memory (MB) field, skip to the next step.

If you changed any of the Buffer Pool Size fields, enter the following command to update the affected databases configured on your server:

fn_util update<xxxx>

where **<xxxx>** is the MKF database (for example, **trans**, **perm**, or **sec**) whose buffer pool size was modified.

Repeat this step for each MKF database whose buffer pool size was increased.

4 Stop all FileNet processes by entering the following command:

killfnsw -DAy

5 Enter the following command to bring up the FileNet software:

initfnsw start

Note any errors when restarting the FileNet software after a configuration change. In many cases, the size of the recovery logs also need to be increased in proportion to the new buffer pool sizes.

Printing Tab

The Printing tab enables you to view and modify parameters for print services and specific printers (or Fax machines). These parameters are located on two subtabs.

Print Services Subtab

Print Services is a subtab of the Printing tab. This subtab contains the following fields:

Server Name

This field displays the host name for the server. The name is created initially by the setup program when the software is installed. This field is locked in this spreadsheet and cannot be

edited here; to change the server name, go to the Server Attributes tab.

In Library Timeout

This field displays the number of seconds the software will wait before a timeout error occurs when requesting a document from storage media resident in the storage library. The default value for this field is 3600.

Out of Library Timeout

This field displays the number of seconds the software will wait before a timeout error occurs when requesting a document from storage media not resident in the storage library. The default value for this field is 3600.

Cache Elasticity

When the print cache fills beyond the fullness threshold, data is sent from the cache to the printer. You can set values in this field from 0 through 6. The value set is used to calculate the fullness threshold. The formula for the calculation is:

fullness threshold = 100% - ($value * 10)

For example, if this field has a value of 2, the print cache will stop accepting new print requests and will start printing its contents when the cache is 80% full.

fullness threshold = 100% - (2 * 10) = 80%

Average Bytes Per Page

This value is used to calculate the number of pages that will fit into the print cache. The default is 50,000 bytes. If the value is set too low, you will fail to utilize all available print cache. If the value is set to high, you could overflow the cache. In this case,

print services deletes objects from the cache before they are printed, then remigrates them from storage media for printing at a later time.

Bad Request Keep Alive (Hr)

This variable will allow you to control how long an unserviced print request, which terminated with an error, will remain in the MKF database. This is useful to help determine which request may have failed, and why. If the print request succeeds, the request is deleted from the MKF database. The default is 24 hours, but you can edit this field to specify a longer or shorter period.

Printers Subtab

Printers is a subtab of the Printing tab. This subtab contains the following fields:

Server Name

This field displays the host Name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Printer Ordinal

This field displays the printer number of each printer installed on the server listed in the column to the left. The number indicates the order in which each printer was added to the server.

The first printer added is assigned a printer number of 0, the second printer is assigned a printer number of 1, and so on.

Printer Name

This field displays the Network Clearinghouse object name for

the printer. The printer name is initially created when the printer is first added to the system under the Procedures tab.

Printer Type

The printer type field lists the currently supported printers. If you want to change the printer type, use the pulldown arrow to select a different printer.

WorkFlo/Print Address

This is the unique TCP internet address assigned to the print station.

The TCP address must be a four part number. Each part must be between 0 and 255 and separated by a period. (e.g., 135.0.20.91)

Cache Allocation

This number, between 1 and 10, specifies the relative cache allocation for each printer you have configured. For example, if you have two printers, setting the cache allocation for both printers to 5 will provide equal cache allocation. If one printer is more active than the other, you might set the cache allocation for the more active printer to 6, and set the cache allocation for the other printer to 4. Click on the arrow in this field to incre- ment the number. To lower the number, continue clicking on the arrow to cycle back to 1.

Paper Sizes

This field displays the current paper sizes defined for this printer. Use the pulldown arrow to display theSelect Supported Paper Sizes dialog box. Select the paper sizes you want from this dialog box; then click the OK Button to complete the config-

uration, Backup to return to the previous window, or Cancel to exit from the procedure.

Paper Trays

This field identifies the number of paper trays in the printer.

Eject Trays

This field lists the trays that are designated as eject trays. Use the pulldown list to select a new eject tray or modify this field.

Fax Modes

This field indicates the Fax Mode or resolution for the fax machine listed. The Fax Mode is set when the fax machine is added to your system, but you can edit the Fax Mode here as well. You can select "coarse," "fine," or both.

Maximum Number of Requests

This field contains the number of maximum print requests that can be supported by the printer. The default is 32.

Device Security

Use this field to configure device security, or access, to the printer. Mark the box Yes or No to activate or deactivate device security.

Yes - The printer is treated like a user. It must have read access for a document before it is allowed to print it.

No - No access required. The printer can print anything.

Relational Databases Tab

The Relational Databases tab enables you to view and modify parameters and objects for Oracle, DB2, or Microsoft SQL Server relational databases.

Changes to parameters only take effect after a database has been installed. If the database is already installed, changes take effect only after the database is restarted.

Note   Microsoft SQL Server is supported only on the Windows platform.

Oracle Relational Database

The Oracle Relational Database tab is a subtab under Relational Databases. It contains the following fields:

Server Name

This field, which cannot be edited, displays the server's host name. The name is created by the setup program when the software is installed. To change the server name, go to the Server Attributes tab.

Home ID

This field displays an integer value that represents the instance of Oracle in use. Usually, this value is 0 (zero), but it might be a higher number if more than one instance of Oracle is installed on the server.

Version Number

This field lists the version of Oracle currently in use.

Database Files

(Filenet-controlled Oracle only)

The maximum number of Oracle database files that can be opened at runtime. The system assigned default value is 20. Changes to this value will not take effect until the database is restarted.

Sort Area Size (Bytes)

(Filenet-controlled Oracle only)
The amount of Program Global Area memory (in bytes) that will be used for sorting rows of data. The system assigned default value is 131,072 bytes. The PGA memory is released once the sorted rows of data are written to disk.

Log Checkpoint Interval

(Filenet-controlled Oracle only)
The number of newly filled redo log blocks (in System Global Area memory), needed to trigger a checkpoint. Once a check-point has been reached, the DBWR (Database writer) process writes all the modified blocks to the disk's redo log. The default is 50,000 blocks.

Index DB Table Size

(Filenet-controlled Oracle only)
The size of the FileNet tables space for the Index database. Four options are available under the pulldown arrow:

Small - Select small if fewer than 1 million documents will be entered into the system.

Medium - Select medium if 1 to 10 million documents will be entered into the system. This is the system default value.

Large - Select large if more than 10 million documents will be entered into the system.

Demo - Select demo if fewer than 10 thousand documents will be entered into the system.

License Maximum Sessions

(Filenet-controlled Oracle only)

The maximum number of user sessions that can occur simultaneously. Once the maximum is reached, only users with Restricted Session privilege can connect. The system assigned default value is 64. For runtime users, License Maximum Sessions is derived from the SLAC Oracle license.

Log Archive Start

This field determines if archiving of the redo log will be automatic. This field must be checked Yes if you are doing Online Backups of the FileNet software. The system default is No.

Click on this field to toggle between Yes and No.

Note: Archive logging must already be enabled in Oracle.

Archive Log Destination

The path name where archived redo logs are kept. The path includes the base name of the archive log file.

If you checked *Yes* in the Log Archive Start field on this tab, you must provide an Archive Log Destination path.

Number of Rollback Segments

(Filenet-controlled Oracle only)

Rollback segments are a dedicated part of an Oracle database that contains database activity information. The information is used to return the database a previous state in case a database action fails.

Oracle requires that there is at least one rollback segment. The default is 4.

To change the number of rollback segments, highlight the current entry and type in a new value.

Additional rollback segments can be added to the oracle_tr dataset (temporary/rollback data). To create new rollback segments, run "fn_util updaterdb."

Note   If the number of rollback segments in the database is less than the number listed in the init.ora file, the database will not start.

Languages

(Filenet-controlled Oracle only)
This field contains the language parameter used to set the Oracle nls_territory parameter in the init.ora file. The territory is the name of the country associated with the language selected.

Select the language you want from the drop-down list.

Character Sets

(Filenet-controlled Oracle only)
Several character sets are available for this field. Select the character set you want from the drop-down list.

Note   The character set is established when initializing the Oracle database. Changing the character set thereafter has no effect on the database.

DB2 Relational Database

The DB2 Relational Database tab is a subtab under Relational Databases. It contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited.

Version Number

This field lists the version of DB2 currently in use.

Password Expiration Policy

This field lists the number of days that the f_sw, f_maint, f_sqi, and f_open passwords remain in effect before they expire. The default value is 60 (days). To change the default, enter a new value in this field.

Note   A blank field is not permitted, and a value of 0 is equivalent to "Never Expires."

To change the passwords themselves, before they expire, use the Application Executive (Xapex).

CAUTION   Once the passwords expire, the IS software (including the Application Executive) will not run and you will have to update the passwords by running Xdb2connect -r from a command line.

Notify Administrator

This field lists the number of days prior to password expiration that the administrator will be reminded to update the password.

The default value is 14 days before the password expires. To change the default, enter a new value in this field.

Note   This value must always be less than or equal to the password expiration value. A blank field is not permitted, and a value of 0 would mean notification the day the password expires.

MSSQL Relational Database

MSSQL is supported on Windows servers only. The MSSQL Relational Database tab is a subtab under Relational Databases. It contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Remote Server Name

If you have an instance on a Remote Server, enter the name of the server where MSSQL resides in this field. If you have a local server, leave this field blank..

Instance Name

Beginning with MSSQL 2000, Image Services supports multiple instances: the "default instance" and other "named instances."

This field allows you to type in the name of the MSSQL instance that you want the IS to use. Any valid string, up to 127 characters, will be accepted.

The default instance name is "MSSQLServer." If the "instance name" field is blank, or contains the name "MSSQLServer," the default instance will be used. If the field contains any other name, that named instance will be used.

The instance can reside on the Remote Server, or on the local server. If the instance is on a Remote Server, the name of that server must be entered in the "Remote Server" field. If the instance is on a local server, the "Remote Server" field should be left blank.

User Connection

This variable sets the maximum number of user connections that can be connected to the Microsoft SQL Server at the same time. It does not refer to the maximum number of processes; that number depends not only on this value but also on other system activity. The default is 20.

Memory

This variable sets the size of memory, in 2K units, the Microsoft SQL Server allocates from the operating system. The default value of memory is 4096.

The more memory available, the more resources Microsoft SQL Server has for internal buffers and caches, reducing the number of times the server has to read data from disk for static information or compiled procedure plans.

There is no performance penalty for configuring Microsoft SQL Server to use the maximum memory available to it on your system. However, you must be sure to assess other memory needs on your system, or Microsoft SQL Server may not be able to acquire enough memory to boot.

Procedure Cache (%)

This variable gives the percentage of memory allocated to the procedure cache after Microsoft SQL Server's memory needs are met. The default is 30 percent.

Microsoft SQL Server's memory needs are the sum of memory necessary for locks, user connections, the code itself (which varies slightly from release to release), and other resources. The remaining memory is divided between the procedure cache and the data cache according to the percentage given by this configuration variable.

Open Objects

This variable sets the maximum number of database objects that can be open at one time on Microsoft SQL Server.

The default value is 500. If this number proves insufficient, Microsoft SQL Server displays a message.

Setting the number of open database objects higher does not have a significant impact on performance or storage requirements. Therefore, you may have to increase this value if Microsoft SQL Server displays a message saying that you've exceeded the allowable number of open database objects.

Recovery Interval

The recovery interval sets the maximum number of minutes per database that Microsoft SQL Server should use to complete its recovery procedures in case of a system failure. The default value is 5 (minutes per database).

The MSSQLServer uses this number and the amount of activity on each database to decide when to checkpoint each database. When Microsoft SQL Server checkpoints a database, it

writes all modified pages (data pages which have been changed by data modification commands) to the disk. The checkpoint also performs a few other "housekeeping" tasks.

You may need to change the recovery interval if your application and its usage change.

Default Language

This is the language that is used to display system messages. The default (us_english) is set at installation time. Use the pull-down menu to change the default setting.

RDB Objects

The RDB Objects tab is a subtab under Relational Databases. An RDB Object can be any FileNet object in an RDBMS, such as a specific table, index, rollback segment, or WorkFlo queue. These objects can be defined by users or other FileNet subsystems.

The RDB Object tab contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Object Name

This field displays the RDB Object Name.

Location

This field displays the Location of the object named in the Object name field. When an object is defined, an RDBMS location must be specified. The location is an existing Oracle

tablespace name, DB2 tablespace name, or an MSSQL Server filegroup name.

To change the location, highlight the text in this field and type in a new location.

Server Appl. Services Tab

This tab enables you to view and modify parameters for Server Application Services. Parameters for these application services are located on the subtabs listed below.

Scheduling Subtab

Scheduling is a subtab of the Server Application Services tab. The parameters on this subtab help the scheduler process determine when to unload media currently in the storage drive in favor of media that have write requests waiting for them.

This subtab contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

RSVP/INFO Trigger

The RSVP/INFO Trigger is used to launch a user-provided program or script that will automatically generate email notification messages to inform the System Administrator that an RSVP or INFO event has occurred. The System Administrator can then respond to the RSVP or INFO message. This feature is

optional. If you choose not to implement this feature, there will be no change to the behavior of your system.

To configure this optional feature, enter the full path to the script or program in this field.

Write Delay Values

The following three parameters help prevent excessive disk swapping due to write requests for several different media. These parameters should minimize the number of media swaps for write requests and optimize the processing of such requests.

For example, if any of the three values below are exceeded and a write request for new media is made, the current disk is swapped for the new disk. If medium or high priority read requests for a new disk are pending, these write values do not apply. The current disk will be swapped regardless of the Write Delay Values. The current disk will also be swapped for a new disk if the retrieval cache is more than 85 percent full.

Write Delay Requests

Number of write requests that can accumulate for all out-of-drive media. If this number is exceeded, a new disk is mounted. The default is 50.

Write Delay Time (sec)

Number of seconds that the scheduler will wait for a write request for the disk currently in the drive. If this number of seconds is exceeded, the disk can be replaced with another (out-of-drive) disk that has write requests. The default is 60.

Mount Delay Time (sec)

This field contains the number of seconds that the scheduler can keep a disk in the drive. If this number of seconds is

exceeded, the disk can be replaced with another (out-of-drive) disk that has write requests. The default is 600.

Prefetch Delay Values

The following two parameters help prevent excessive disk swapping due to prefetch requests for several different media. These parameters should minimize the number of optical disk swaps for prefetch requests and optimize the processing of such requests. Prefetch requests are the lowest priority of requests. They will only cause a swap to occur if there are no higher priority requests (read or write) pending.

Prefetch Delay Requests

Number of prefetch delay requests that can accumulate for any out-of-drive disk. If this number is exceeded, (and there are no higher priority requests pending) a new disk is mounted. The default is 0.

Prefetch Delay Time (sec)

Number of seconds that the scheduler will wait for a new prefetch request for the disk currently in the drive. If this number of seconds is exceeded, the disk will be replaced with another (out-of-drive) disk that has prefetch requests. The default is 0.

Minimum Transfers

This field contains the number of operations (document reads, writes, etc.) that should be done before the disk is swapped. This value applies only if the disk in the drive has the same priority as the (out-of-drive) disk requested. The value should be as large as your retrieval cache can support. You should consider the average document size (as well as the retrieval cache size) when you set this value. The default is 50.

Critical Read Time (sec)

> This field contains the number of seconds a document request can wait before it becomes critical to mount the disk containing the requested document. The scheduler uses the critical time to determine which disks to mount, as follows:

- If the wait times for the outstanding document requests have not exceeded the critical time, the scheduler selects the storage media surface with the most requests.

- If the wait times for the outstanding document requests have exceeded the critical time, the scheduler selects the storage media surface with the oldest request.

> This formula achieves a balance between optimum turn-around (mount disks with old requests) and optimum throughput (mount disks with many requests). The default is 60.

I/O Active Time (sec)

> This field contains the number of seconds in time span used by the scheduler to keep track of read/write requests for each disk in a drive.

> The scheduler knows how many read/write requests were made for these media in the last time span. The disk with the fewest requests (the least active disk), is the disk most likely to be replaced with another (out-of-drive) disk that has write requests.

> This parameter is primarily for systems with low I/O activity. The default is 1800.

Spindown Delay (sec)

> This field contains the number of seconds a disk can be idle before it is spun-down (a preliminary activity to removing a disk

from the drive). If this number is exceeded, the disk is spun-down. The default value is 6 seconds.

This parameter applies only to OSAR-GT devices. The parameter helps balance the time required to move the gripper and the time required to spin-down a disk.

If the gripper is frequently waiting for the disk to spin down before removing it, you may want to lower this parameter value. A lower value increases the probability that the disk will begin to spin down before the gripper moves to the drive; therefore, the gripper will not need to wait as long to remove the disk. The default is 6.

Read Look Ahead Count

The data transfer process looks ahead for a specified number of read requests in order to minimize seek time on storage media. Of these read requests, the process executes the one which is at the closest sector address to the read request just executed. The default value for this field is 10.

OSAR Server ID

This read-only field contains the system assigned ID number for the server providing Osar Services. The default is 2.

Cache Duration Subtab

This subtab of the Server Application Services tab contains four fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Prefetch Cache Duration (sec)

This field contains the number of seconds that a prefetched document is guaranteed to remain in the pagecache. It can be deleted from the cache after this, unless it is requested again. The default is 300.

Migrate Cache Duration (sec)

This field contains the number of seconds that a retrieved document is guaranteed to remain in the pagecache. It can be deleted from the cache after this, unless it is requested again. The default is 300.

Refresh Cache Duration (sec)

For a document that was requested again (requested while already in the cache), this field contains the number of seconds the document is guaranteed to remain in the cache. The default is 120.

Batch Subtab

This subtab of the Server Application Services tab contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Target Image Services Domain Name

This field specifies which system (target system) will receive the committed batches. Use the pull-down menu to select only compatible systems as remote target systems. Compatible sys-

tems have no overlapping document IDs or media IDs. Select 0 to specify the local system as the target system.

Fast Batch Committal

Fast batch committal commits groups or batches of documents together for more efficient committal. It is a performance tuning item.

Note If fast batch committal is enabled, you cannot use cluster indexes. See the ***System Administrator's Handbook*** for more details on fast batch committal and clustering.

If this option is turned on, fast batch objects are automatically deleted from cache after they have been committed to an optical device. To retrieve an image that has just been committed, the image has to be migrated back to cache.

In addition, drives doing mostly fast batch committal should favor writes over reads.

When the Fast Batch Committal option is turned on, it is not necessary to turn on the Delete On Write option. The objects will be deleted from cache regardless.

However, if the Fast Batch Breakup option is also turned on, it will take precedence over both the Fast Batch Committal option and the Delete On Write option, and the objects will remain in cache.

To enable Fast Batch Committal, check this box Yes.

Cache Subtab

This subtab of the Server Application Services tab contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Cache ID

This read-only field contains the system assigned ID number for the cache.

Cache Type

This read-only field contains the Cache Type, which can be: Batch, Retrieval, System Print, or Application Print.

Note   Revise, Fill-in, and Folder View caches are NOT SUPPORTED in this release.

Batch Cache	Batch cache holds document entry batches between the time the batch is created and the time it is ready to be committed. When the batch is ready to be committed, it is moved to the retrieval cache.
Retrieval Cache	Retrieval cache holds documents for retrieval by workstations and documents waiting to be written to storage media.
	When Batch Entry services is ready to commit a document, it moves the document from the batch cache to the retrieval cache. The document tables are then updated and the document is written to storage media.

System Print Cache	System Print cache holds printable objects, such as Apex screens, text files, COLD documents, and bitmap images, waiting to be printed. System print cache is used only for image migration.
Application Print Cache	Application Print cache holds printable objects, defined by applications, waiting to be printed. WorkFlo/Fax uses application print cache for its mailboxes. The application print cache is configured automatically, and there must be exactly one application print cache per Print Service.
Revise Cache	Revise cache is NOT SUPPORTED in this release.
Fill-in Cache	Fill-in cache is NOT SUPPORTED in this release.
Folder View Cache	Folder View cache is NOT SUPPORTED in this release.

Minimum Allocation (%)

This field contains the minimum percentage of the total physical cache space that should be reserved for the logical cache you are setting up. This amount of space is always available for the cache, which also means that no other caches can use this space.

The minimum and maximum percentages assigned to the retrieval cache must be equal.

Maximum Allocation (%)

This field contains the maximum percentage of the total physical cache space that should be reserved for the logical cache you are setting up. The cache cannot use any more space than you allocate with this value.

You must not specify 0 as the maximum percentage for a cache, even if you do not plan to use the cache. (This require-ment is due to internal software dependencies.)

Note   For Retrieval cache, the minimum and maximum percentages must be equal. The Maximum Allocation (%) automatically assumes the same value as the Minimum Allocation (%) field.

The Maximum Allocation (%) field for Revrieval cache cannot be edited.

The media sizing tool can help you determine the maximum percentage for each cache. The number in the Percent column can be rounded up and used as the maximum percentage.

The total of maximum percentages can be more than 100 per-cent. However, this configuration may create situations in which a cache with a 0 minimum percentage may not have any space available for it.

Locked Threshold (%)

This field contains the percentage of cache (allocated to each cache type) that can be filled with locked documents before the system will stop accepting uncommitted images.

If the number of locked documents exceeds this percentage, you will not be able to scan documents. You will not be able to resume scanning until enough images are written to media to put the number below the locked percentage.

Example: If Batch cache is set with a maximum allocation of 50% out of a total cache of 10GB, then a 90% threshold would be triggered when the Batch cache reached 4.5GB.
(5GB x .90 = 4.5GB)

You should accept the default unless you use overnight prefetching. If prefetching is done, the percentage must be smaller than the Cache threshold write percentage by either 5000 sectors or 10 percent of the retrieval cache, whichever is less.

Daemon Threshold (%)

This field contains the percentage of cache (allocated to each cache type) that can be filled with both unlocked and locked documents before the cache daemon will begin deleting unlocked images.

If the number of unlocked and locked images exceeds this percentage, unlocked images will be deleted until the combined number of images no longer exceeds the percentage.

Example: If Batch cache is set with a maximum allocation of 50% out of a total cache of 10GB, then a 90% threshold would be triggered when the Batch cache reached 4.5GB.
(5GB x .90 = 4.5GB)

You should accept the default unless you use overnight prefetching. If prefetching is done, the percentage must be less than the Cache threshold write percentage by either 10,000 sectors or 10 percent of the retrieval cache, whichever is less.

Write Threshold (%)

This field contains the percentage of cache (allocated to each cache type) that can be filled with both unlocked and locked documents before storage media writes are given a higher priority than storage media reads. Usually, reads have a higher priority than writes.

Example: If Batch cache is set with a maximum allocation of 50% out of a total cache of 10GB, then a 90% threshold would

be triggered when the Batch cache reached 4.5GB.
(5GB x .90 = 4.5GB)

You should accept the default unless you use prefetching. If
prefetching is done, the percentage must be less than the total
retrieval cache size (for example, 100 percent) by either 5000
sectors or 10 percent of the retrieval cache, whichever is less.

ICR Subtab

ICR (Intelligent Character Recognition) is NOT SUPPORTED in
this release.

Server Attributes Tab

The Server Attributes tab contains the following fields:

Domain Name

This field displays the Network Clearinghouse domain name.
The name is initially created by the setup program when the
software is installed. This field is locked and cannot be edited.
Go to the System Attributes tab to edit the domain name.

Server Number

The Server Number is the sequential number automatically
assigned to each server when the system is configured. For
example, the server number of a combined server is 1. This
field cannot be edited.

Server Name

The field contains the name of the server whose configuration
database you are currently editing.

Server Type

The Server type in this read-only field indicates whether this server is a:

Combined (Root/Index/OSAR) server

Dual (Root/Index) server

Root/OSAR server

Index server

DocLocator server

OSAR server

Application Server

Platform

The Platform listed in this field indentifies the type of computer hardware on which the FileNet Image Services is installed. Platform types include:

IBM RS/6000

HP Integrity

HP 9000/800

Windows Server

Sun (Solaris)

Use the pulldown menu to modify the platform type.

Services Tab

The Services tab lists the services installed on your system server. This tab contains the following read only information:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Batch Services

If this field is checked Yes, it indicates this server is responsible for Batch Services. Go to the System Defaults tab to select a different server.

Cache Services

If this field is checked Yes, it indicates this server is responsible for Cache Services. This service is automatically assigned to any server providing Batch or Print Services.

Print Services

If this field is checked Yes, it indicates this server is responsible for Print Services. Go to the System Defaults tab to select a different server.

WQS Services

If this field is checked Yes, it indicates this server is responsible for WorkFlo Queue Services. Go to the System Defaults tab to select a different server.

SQL Services

If this field is checked Yes, it indicates this server is responsible

for Structured Query Language Services. Go to the System Defaults tab to select a different server.

ICR Services

ICR (Intelligent Character Recognition) is NOT SUPPORTED in this release.

WMS Service

If this field is checked Yes, it indicates that this server has Workflo Management Services installed.

System Appl. Services Tab

The System Application Services tab allows you to view and modify system-wide application services parameters located on two subtabs.

Document Services Subtab

This subtab of the System Application Services tab contains seven fields:

Domain Name

This field contains the name of your system's domain.

Document Number (low value)

This is the first number in the range of document numbers to be assigned on your system. This value must be 100,000 or greater. The start of the document number range must not overlap the ranges of other compatible systems.

Document Number (high value)

This is the last number in the range of document numbers to be assigned on your system. This value must be greater than the Document number low value, and be 3,999,999,999 or less.

The end of the document number range must not overlap the ranges of other compatible systems.

Surface ID (low value)

This is the first number in the range of surface IDs assigned on the system. This value must be 3000 or greater. The start of the surface ID range must not overlap the ranges of other compatible systems.

Surface ID (high value)

This is the last number in the range of surface IDs to be assigned on your system. This value must be greater than the Surface ID low value, and also must be 4,294,967,290 or less. The end of the surface ID range must not overlap the ranges of other compatible systems.

Pages Per Retrieval

This field contains the number of pages retrieved with each document display request. The maximum allowable value is 300. A smaller number improves system performance. A larger number provides more pages for a user to "browse" through with good response time. The default is 8.

Document Batch Buffer Size

This field specifies the size of the buffer used for fast batch committal. The default value, 80 KB, multiplied times 1024 is the buffer size used. The maximum value is 511 KB. This value applies to all servers with Document Services.

Other Services Subtab

This subtab of the System Application Services tab contains the following fields:

Domain Name

> This field contains the name of your system's domain. The name entered here is used as the "domain" part of your system's NCH name.

Delete on Write

> This field specifies whether documents are deleted from the page cache after they are written to storage disk.
>
> Click on this field to toggle the setting to Yes if you want documents to be deleted after they are written to disk.

Note This option is ignored if the "Fast Batch Breakup" option is turned on.

Fast Batch Breakup

> This option provides Image Services with the capability to breakup fast batch objects during the committal process, leaving retrievable document pages in the storage library server's cache. This allows an image to be retrieved directly from cache after it's been committed, without having to retrieve it from the optical device or MSAR library. If this option is used, the need to prefetch these documents into cache is eliminated.
>
> If either the Delete On Write option or the Fast Batch Committal option is also turned on, they are ignored and the objects will remain in cache regardless.
>
> If you want to use the Fast Batch Breakup option, turn this feature **ON** by clicking in this field. The default is **Off**.

Note If you have a Cache-only system, this option is ignored.

Optical Disk Read After Write Count

This field specifies how often to do read after write checking for optical disk writes. If the value is zero (default), then no read after write checking is done. If the value is 1, then every write is checked by a read after write. If the value is "n," then every "n"th optical disk write is checked with a read after write.

The highest-performance, lowest-reliability setting is 0 (never do a read after write). The lowest-performance, highest-reliability setting is 1 (always do a read after write). After that, as the number gets higher, the performance goes up and the reliability goes down (because you are doing read-after-write less often).

Click in this box to type in a new value.

Checksum

Checksumming is used to verify image integrity. If you suspect hardware problems, checksumming can assist you in pinpointing the difficulty. Checksumming has an impact on system performance, but it provides you with a valuable troubleshooting tool. The available settings are: None, Minimum (the default setting), or Maximum. To modify this setting, use the the pulldown menu.

None

If checksumming is set to None, any checksums (if present) will only be validated by the server during normal committal, at the point where they are written to optical disk. When set to None, some remote FileNet applications, such as Document Entry, will not generate documents with checksums.

Minimum

This is the default setting for this field. The Minimum default

value only applies to new installations. Systems upgraded from a previous release maintain their previous checksum setting.

Minimum checksumming (or end-point checksumming) causes a checksum to be calculated as soon as the image data is written to magnetic disk. At that point, the checksum is used to verify that the image is intact. When checksumming is set to Minimum (or Maximum), some remote FileNet applications will automatically generate checksums for documents.

Maximum
Maximum checksumming causes the checksum to be recalculated and checked each time the image data is moved anywhere in the system: within memory, to and from magnetic disk, or to and from storage media.

With maximum checksumming, additional verification is done to insure that a corrupted image in cache is detected before it is written to optical. (The images written to cache may have been corrupted because of magnetic disk faults.) After the image is written to cache, it is read into a temporary buffer and compared to the original cache image.

Note   When checksumming is set to Maximum (or Minimum), some remote FileNet applications will automatically generate checksums for documents.

Cataloging Enabled
This field specifies whether documents are allowed to be entered into the index database. The default is Yes. Click this box to change the entry. If you mark this box (no), no entries are made to the index database and no index cataloging is performed during the committal phase. In either case, the docu-

ment indices are still inserted into the WorkFlo queue if the document class has a WorkFlo queue associated with it.

Dist. Queue Mode

This field specifies whether other items besides the Document ID will be inserted into the WorkFlo queue when committing documents. Click this box to change the entry.

Yes

If you specify Yes, all user indices whose names and types match fields in the queue will be passed to the respective field, and the document number will be inserted into the first WorkFlo queue field of type document.

No

If you specify No, only the Document ID number to be inserted into the first document type field of the WorkFlo queue is allowed.

Collect Statistics

This field specifies whether Quick Logging statistics will be maintained. Click this box to change the entry. If you check this box (yes), statistics for committals, retrievals, and deletions during system image processing are kept in the /fnsw/local/logs/qlogs/1 directory.

Surface Activity Tracking

This field allows you to monitor optical or MSAR activity on a per surface/per day basis. Click this box to change the entry. This information can be used to help you predict system needs, and which surfaces are good candidates for off-line storage. Surface activity reports are run via the DOC_tool utility.

Export Annotations

This field specifies whether annotations will be copied to a P8

Content Engine system in a Content Federation Services for Image Services environment. The default setting is No.

Yes

If you specify Yes, annotations will be copied to the Content Engine server whenever they are created or updated. Set this field to Yes when the Content Engine system is running P8 4.0 or higher.

No

If you specify No, any new annotations or changes to existing annotations will not be copied to the Content Engine system.

System Defaults Tab

The information on this tab indicates the system-wide services and default devices that have been defined for this FileNet Image Services system.

Domain Name

This field contains the name of your system's domain. The name shown here is used as the "domain" part of your system's NCH name.

Batch Services

This field indicates the name of the server where Batch Services are located on the system. Use the pulldown list to select the server where Batch Services are located.

If no server name is shown, the service is not active anywhere on the system.

Print Services

This field indicates the name of the server where Print Services

are located on the system. If no server name is shown, the service is not active anywhere on the system.

WQS Services

This field indicates the name of the server where WQS (WorkFlo Queue) Services are located on the system.

If no server name is shown, the service is not active anywhere on the system.

SQL Services

Indicates whether SQL (Structured Query Language) Services are active on the system. If no server name is shown, the service is not active anywhere on the system.

ICR Services

ICR (Intelligent Character Recognition) is NOT SUPPORTED in this release.

Printer Name

This field contains the name of the system default printer.

To change the system default printer, use the pulldown list to select another printer as the default. To add a new printer, use the Procedures tab. To modify the characteristics of an existing printer, use the Printing/Printers tab.

If no printer names are shown in the pulldown list, no printers have been installed on the system.

Tape Drive Name

This field contains the name of the system default tape drive. The tape drive shown may or may not be the same as the server default tape drive shown on the Tape Drives tab.

To change the system default tape drive, use the pulldown list to select another tape drive as the system default. To add a new tape drive, use the Procedures tab. To modify the characteristics of an existing printer, use the Tape Drives tab.

If no tape drive names are shown in the pulldown list, no tape drives have been installed on the system.

Revise Cache

Revise cache is NOT SUPPORTED in this release.

System Attributes Tab

The System Attributes tab contains the following fields:

Domain Name

This field contains the name of your system's domain. The name shown here is used as the "domain" part of your system's NCH name.

Software Release

This field contains the Software Release number of the current FileNet Image Services software installed on this server. This field cannot be edited.

System Type

This field contains the system type. It is always shown as IS. This field cannot be edited.

System Serial Number

Your system serial number is provided to you by FileNet.

For UNIX systems, this number is based on the unique processor ID of the server.

For Windows systems, this number is based on the network adapter address. You can edit this field.

Note If a new server is installed or the processor ID or network adapter is replaced, the system serial number changes and you must obtain a new SLAC key.

Default Character Set

This field contains the default character set used for indexing.

Former Character Set

If the system has been converted from one character set to another, the Former Character Set indicates the character set previously used for indexing. Normally, the Former Character Set and the Default Character Set are the same.

Network Protocol

This field contains the network protocol the system currently uses. Only TCP/IP is currently supported.

Protocol Preference

Since only one network protocol is supported, and listed in the Network Protocols field (to the left), that protocol is also listed here as the preferred protocol. This field can only contain the TCP/IP network protocol.

Tape Drives Tab

The fields on this tab contain the parameters for tape drives that are installed on your system. If no tape drives are installed, all fields will be blank.

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

Tape Drive Type

This field indicates the type of tape drive installed. Use the pull-down menu arrow to change the type of tape drive.

Device Address

This field is not currently used.

Device Name

This is the unique Network Clearinghouse object name for your tape drive (for example, Tape1). The name is initially created when the tape drive is first added to the system under the Procedures tab.

Default Tape Drive

This field is used to set a tape drive as the default unit for its associated Server Name. Set the field to Yes for the tape drive you want as the default unit. Click on the field to change the value between Yes and No.

If only one tape drive is configured on your system, it is automatically designated as the default drive and you will not be able to change this field.

Only one drive can be designated as your default server tape drive. The name of the system default tape drive may be different and can be viewed (or modified) under the System Defaults tab.

WorkFlo Mgmt. Services Tab

The WorkFlo Mgmt. Services tab enables you to view and modify parameters of WorkFlo Management Services and Servers that are currently configured on your system.

There are two subtabs, VW Servers and VWServices.

Note Throughout the System Configuration Editor interface, workflow servers and services are referred to as VW Servers and VWServices.

VW Servers Subtab

VW Servers is a subtab of the WorkFlo Mgmt. Services tab.

A VW Server may be added, or deleted, using the appropriate procedure on the Procedures tab. This subtab contains the following fields:

Server Name

This field displays the host name for the server. The name is initially created by the setup program when the software is installed. This field is locked and cannot be edited. To change the server name, go to the Server Attributes tab.

VWService Number

This field displays the VWService number associated with the named WorkFlo Server. This number uniquely identifies the

VWService throughout an IS domain. The VWService number is automatically assigned for you, starting with 0.

VW Server Ordinal

This field identifies the ordinal number assigned to the named WorkFlo Server. The ordinal numbers are unique within a given VWService and are automatically assigned for you, starting with 0.

Note   A VWService may have more than one server associated with it.

VWServices Subtab

VWServices is a subtab of the WorkFlo Mgmt. Services tab.

An IS domain may have one or more VWServices and each VWService may have one or more WorkFlo servers. A VWService may be added, deleted, or modified by using the appropriate procedure on the Procedures tab. This subtab contains the following fields:

Domain Name

This is the Network Clearinghouse domain name where your system is located.

VWService Number

This field identifies the VWService number. This number uniquely identifies the VWService throughout an IS domain. The VWService number is automatically assigned for you, starting with 0.

SQL Server Name

This field identifies the server name of the SQL server that is associated with this VWService.

Procedures Tab

When you need to add or delete a configuration item, such as a printer, the Procedures tab is used. For help running any procedure, click the Help button on the Procedures tab.

After you've completed a procedure, your new configuration values can be viewed (or modified) under the appropriate System Configuration Editor tab.

To run a procedure:

1   Click on the procedures tab.

2   Click on the procedure you want to run.

3   Click Run to begin the procedure.

A series of dialog boxes will prompt you for the information needed. When you complete the information in each dialog box:

Click **OK** to continue

Click **Backup** to return to the previous dialog box

Click **Cancel** to exit from the procedure

4   When you have completed making all the changes you want to your configuration, select Exit from the file menu and save your changes.

A new configuration database file is created with the next higher version number.

Note: You must run Task Manager to restart the FileNet software in order for the changes you made to take effect.

List of Procedures

The following is list of procedures contained on the Procedures tab. If you want to find information on a specifc procedure you can go directly to the page number listed, or if you are viewing this document Online simply click on the procedure you want.

Add a Printer

To add a printer or fax machine to your system:

1 Scroll through the list of the currently supported printers and select the type of printer that you want to add.

2 Enter a unique object name for your printer (for example, Printer1).

3   Enter your TCP address for the printer.

The TCP address is a four part number. Each part must be a number between 0 and 255 and separated by a period (for example, 135.0.20.91).

4   After entering the network address:

Click **OK** to continue

Click **Backup** to return to the previous dialog box

Click **Cancel** to exit from the procedure

5   Select paper sizes for the new printer.

This dialog box is used to select paper trays that contain the paper sizes to use. The selection list displays only the supported paper sizes for the type of printer you are adding. Scroll through the list and select the paper sizes you want. Use your mouse button and the Shift or Control key to select more than one tray size.

6   Select the trays to use as eject trays for the new printer. The selection list displays only the supported paper trays for the type of printer you are adding to your system. Default is the only selection available with this release.

7   Select the Fax Mode resolution to use for the fax machine you are adding: Coarse, Fine, or Both.

8   Choose whether to make the new printer or fax machine the default.

Only one printer or fax can be designated as your system default.

Delete a Printer

To delete a printer or fax machine:

1 Scroll through the list of system printers and select the one to delete.

2 If you delete the default printer, you see this message:

System Default Printer is Set to Null

3 If you delete the default printer, you must then go the the System Defaults tab and select another printer to be the system default printer.

Automatically Configure Optical Storage Library

An optical storage library must be installed and powered on for the configuration software to access it. Click Run to automatically configure an optical storage library. When the configuration is complete, you can customize default settings in the appropriate System Configuration Editor tabs.

Note: The auto configuration utility is run on the server physically connected to the storage library device. If the storage library is physically connected to a different server, you must manually configure it.

If a library or drive fails to configure automatically for any reason, use the manual configuration procedure.

The following are limitations to the auto configuration procedure.

• FileNet OSARs cannot be automatically configured because they do not have SCSI arms.

• The Hitachi high capacity, backwards compatible optical drive cannot be automatically configured because its product ID is not

unique. The product ID is the same as the non-backwards compatible drive.

- The FileNet MOSAR cannot be automatically configured because the device is not SCSI II compliant.

Manually Configure Optical Storage Library

This procedure is used to manually configure an optical storage library.

CAUTION   There is a known bug when using the Quit button to cancel this procedure. The Quit button (Cancel feature) does not work from any dialog box after entering the storage library number, and before completing the procedure. Clicking the Quit button results in the storage library being added anyway. If this situation occurs, you must delete the storage library using the Delete a Storage Library procedure.

To manually configure an optical storage library on your system:

1   Select the optical storage library type.

Using the up and down arrows, scroll through the list and select the storage library type that you want to configure.

2   Enter the storage library number.

Select from the list or type in the storage library number.

3   Enter the number of storage drives.

Select from the list or type in the number of storage drives you want to configure.

4   Select the storage drive type.

Using the up and down arrows, scroll through the list and select the storage drive type that you want to configure.

5   Enter the storage drive position.

Using the left and right arrows, scroll through the list and select the storage drive position. The number you select depends on the drive type you selected earlier. For FileNet OSAR libraries, the drive positions must start at 0.

6   Enter the SCSI Address for the drive.

The SCSI address is usually a four part number with a space in between each part. If the drive is connected to an IBM 5712 or 5736 adapter card, the drive's SCSI address must have five parts.

To determine the exact SCSI address of the drive, enter:

    ls /dev/fnsod* (on UNIX servers)

    fndev -v (on Windows servers)

Delete an Optical Storage Library

To delete an optical storage library, use the left and right arrows to enter the number of the optical storage library to remove.

To determine the library's number, find the library's name under the Storage Libraries tab, then convert the name to its correlating number (for example, number 0=Library A, number 1=Library B, etc.).

Configure an MSAR Storage Library

To configre the MSAR library you will have to know the following information:

- Server name (if more than one server is configured)
- MSAR library type
- MSAR library number
- MSAR surface directory
- Number of drives that should favor writes

1 Scroll through the list of servers and select the server to add the storage library to.

Note This dialog box only appears if you have more than one server configured on your system.

2 Select the storage library type.

3 Enter the storage library's number.

To determine the library's number, find the library's name under the Storage Libraries tab, then convert the name to its correlating number (for example, number =Library A, number 1=Library B, etc.).

4 Enter the MSAR surface directory for the MSAR libary. This is the MSAR Creation directory.

Note The MSAR Surface Directory you create here will be listed on the Storage Libraries tab.

5 Enter Number of Drives that Favor Writes.

Since MSAR drives are controlled with software only, an MSAR library always has the maximum number of drives (12) available. To increase performance or for other reasons, you may wish to have fewer active drives. This setting allows you to do that.

To increase throughput in the page cache, some drives can be configured to favor write operations. This allows documents to be unlocked and aged more quickly, thereby increasing performance.

The number you select in this box will apply to the drives begining with drive 1. i.e., If you enter 3, drives 1, 2, and 3 will be set to favor writes.

Enter a number, or accept the default 0, and click **Next** to continue.

Note If you want to change this setting later, use Storage Library Control.

Delete an MSAR Storage Library

To delete an MSAR library you will have to know the library's number. This can be determined by looking up the library's name under the Storage Libraries tab and then converting the name to its correlating number (e.g. number 0 = Library A, number 1 = Library B, etc.).

After you've deleted the library, your configuration values will be removed under the Storage Libraries tab.

1 Click the **Run** button after selecting this procedure on the Procedures tab. The following message appears.

```
If any MSAR surfaces are still in the library, they
should be ejected before deleting the MSAR library. Are
you sure you want to delete the MSAR Library now?
```

2 Click **Yes** to continue.

3 Scroll through the list of servers and select the server to add the storage library to.

Note This dialog box only appears if you have more than one server configured on your system.

4 Enter the storage library's number.

Add an Optical Storage Drive

To add an optical storage drive to an existing library:

1 Enter the optical storage library's number.

To determine the library's number, find the library's name under the Storage Libraries tab, then convert the name to its correlating number (for example, number =Library A, number 1=Library B, etc.).

2 Select the type of storage drive from the scrolling list.

3 Enter the position for the new storage drive in the library.

4 Enter the SCSI address for the new storage drive.

The SCSI address can be determined using fndev tool.

Delete an Optical Storage Drive

To delete an optical storage drive from a library:

1 Scroll through the list of servers and select the server with the storage drive you want to delete.

Note This dialog box only appears if you have more than one server configured on your system.

2 Enter the storage library's number to remove the drive from.

3 Enter the number of the drive to remove.

Use the left and right arrows to enter the physical position number for the storage drive.

Add a Tape Drive

To add a tape drive to your system:

1 Scroll through the list of servers and select the server that you want to configure with a tape drive.

This dialog box only appears if you have more than one server configured on your system.

2 Scroll through the list of the currently supported tape drives and select the appropriate drive type.

3 In the space provided, enter a unique object name for the tape drive you are adding.

4 Choose whether to make the new tape drive the default for the server.

Only one drive can be designated as your default server tape drive.

5 Enter the tape drive device path in the space provided.

This dialog box only appears if you have more than one server config-ured on your system.

6 Choose whether to make the new tape drive the default for the FileNet system.

Only one drive can be designated as your default system tape drive.

Delete a Tape Drive

To delete a tape drive, select the name of the drive you want to delete from the scrolling list.

If the tape drive you delete was the system default tape drive, you must go to the the System Defaults tab and select another tape drive to be the system default. Your system must have a default tape drive.

Set to Site-Control

A site-controlled database is an existing RDBMS (Oracle or MSSQL Server instance) which is not owned or operated by the FileNet IS soft-ware.

To change from a FileNet-controlled database to a site-controlled data-base, select **Set to Site-Control** from the list of procedures and click **Run**.

Set to FileNet-Control

A FileNet-controlled database is a RDBMS (Oracle or MSSQL Server instance) which is owned or operated by the FileNet IS software.

To change from a site-controlled database to a FileNet-controlled database, select **Set to FileNet-Control** from the list of procedures and click **Run**.

Note A password is required to use this procedure. Contact FileNet CSS for further information.

Follow the subsequent prompts as they appear.

Add Relational Database Object

To add a relational database object select/enter the following information in the procedure dialog box.

1 Select the server where you want to add the relational database (RDB) object.

2 Select the RDB object type.

3 Select the RDB record type that you want to use for the RDB that you are adding.

4 Enter the name that you want to use for the RDB object.

5 Enter the location of the RDB object.

Note The combination of object type, record type, name, and location must be unique or you will receive a warning message to re-enter the information.

Delete Relational Database Object

To delete a relational database object:

1 Select the name of the RDB object that you want to delete.

2 Select the type of RDB object that you want to delete and click Next.

After you have the deleted the RDB object, the configuration values will be removed from the Relational Databases tab/RDB Objects subtab.

Set Up Required Oracle Datasets on Root/Index Server

To set up an Oracle dataset on the Root/Index server:

1 Enter the Oracle database path.

The default path for the dataset automatically displays in this dialog box. You can edit the path if needed.

2 Enter the database size in megabytes.

The minimum database size for the new database is displayed in the dialog box. Use the left and right arrows to select the amount of space you want for the dataset.

3 Enter the Oracle redo path.

Oracle requires two redo logs. The default path for the redo log displays automatically in this dialog box. You can edit the path if needed.

4 Enter the redo log size in megabytes.

The minimum size for the redo log is displayed in the dialog box. Use the left and right arrows to select the amount of space you want for the redo log.

Note This procedure only appears if you have an Oracle RDBMS.

Deconfigure Oracle Datasets on Root/Index Server

To deconfigure an Oracle dataset, click the Run button. The Oracle database and redo logs will be automatically removed. Go to the Datasets tab to verify that the Oracle Dataset has been deconfigured.

Note This procedure only appears if you have an Oracle RDBMS.

Set Up Required Microsoft SQL Datasets on Root/Index Server

To set up an Microsoft SQL dataset on the Root/Index server:

1 Enter the Microsoft SQL database path.

The default path for the dataset automatically displays in this dialog box. You can edit the path if needed.

2 Enter the database size in megabytes.

The minimum database size for the new database is displayed in the dialog box. Use the left and right arrows to select the amount of space you want for the dataset.

3 Enter the Microsoft SQL redo path.

Microsoft SQL requires two redo logs. The default path for the redo log displays automatically in this dialog box. You can edit the path if needed.

4 Enter the redo log size in megabytes.

The minimum size for the redo log is displayed in the dialog box. Use the left and right arrows to select the amount of space you want for the redo log.

5 Enter the Microsoft SQL temp space path. The default path for the temp space file displays automatically in this dialog box. You can edit the path if needed.

6 Enter the Microsoft SQL temp space size in megabytes. The minimum size for the temp space file is displayed in the dialog box. Use the left and right arrows to select the amount of space you want for the temp space file.

Note This procedure only appears if you have an MSSQL Server RDBMS.

Deconfigure Microsoft SQL Datasets on Root/Index Server

To deconfigure an Microsoft SQL dataset, click the Run button. The Microsoft SQL database and redo logs will be automatically removed. Go to the Datasets tab to verify that the Microsoft SQL Dataset has been deconfigured.

Note This procedure only appears if you have an MSSQL Server RDBMS.

Create an Optional Relational DB Dataset

To create a new dataset:

1 Select the name of the server where you want to add the dataset from the list displayed.

This dialog box only appears if you have more than one server configured on your system.

2 Select the type of dataset that you want to add from the list displayed.

If you create an Oracle User Database, you must also create Oracle temporary data and rollback segments.

3 Accept the default dataset path or enter a new path.

4 Enter the dataset size in megabytes.

A message displays to remind you to run **fn_util updaterdb** on the system where the new dataset resides. This is necessary to implement the changes you have just made.

5 After you finish your configuration edits, exit the System Configuration Editor and save your changes.

6 As fnsw user, run the **fn_util updaterdb** command at the command prompt.

Delete an Optional Relational DB dataset

To delete an optional relational DB dataset:

1 Scroll through the list of servers and select the server that has the optional relational DB dataset you want to delete.

This dialog box only appears if you have more than one server configured on your system.

2 Scroll through the list of dataset choices and select the dataset you want to delete.

Only one dataset can be selected at a time. If only one dataset is configured on the server you selected above, this dialog box does not appear.

Add an Additional Dataset

To add another dataset:

1 Scroll through the list of servers and select the server where you want to add the additional dataset.

This dialog box only appears if you have more than one server configured on your system.

2 Scroll through the list of available dataset choices and select the appropriate dataset.

3 Accept the displayed default directory for the dataset or edit the pathname.

4 Enter the dataset size in megabytes.

Use the arrows to select the amount of space you need for the dataset. The default size for the new dataset is displayed in the window.

Run "fn_util updaterdb."

A message displays to remind you to run **fn_util updaterdb** on the system housing in order to implement the change you are making.

5 After you finish your configuration edits, exit the System Configuration Editor and save your changes.

6 As **fnsw** user, run the **fn_util updaterdb** command from a command prompt.

Delete an Additional Dataset

1 Scroll through the list of servers and select the server that has the additional dataset you want to delete.

This dialog box only appears if you have more than one server configured on your system.

2 Scroll through the list of dataset choices and select the dataset you want to delete.

Only one dataset can be selected at a time. If only one dataset is configured on the server you selected above, this dialog box does not appear.

Add a Peer System

Peer systems are used to represent both the remote and compatible system concepts. A remote system is recognized by its internet address, is accessable via a network, and is usually a recipient for data from a remote entry station (which has no storage library of its own).

A compatible system is set up with a unique range of document and surface IDs. This allows an easy import of documents and media from one system to the compatible system. Compatible systems are usually not on the network and are recognized by their system serial number.

To add a the new peer system:

1 Enter the Network Clearinghouse domain name of the remote or compatible system.

2 Enter the system serial number.

This is the serial number of the remote or compatible system (mandatory if you are adding a compatible system).

3 Enter the network address for this system.

This is the four part internet address (for example, 135.0.20.91) of the remote or compatible system (mandatory if you are adding a remote system).

Remove a Peer System

To remove a peer system from your configuration, scroll through the list of peer systems, select the system you want to remove, and click OK.

If only one peer system is configured on your system, this dialog box does not appear.

Add an Optional Logical Cache

To add an optional logical cache to your system, scroll through the list and select the logical cache type that you want to add.

If you select a cache type that already exists on the server, a notification message displays. Click the OK button and select another cache type.

Delete an Optional Logical Cache

This procedure is used to delete an optional logical cache to your system. Click the Run button to continue with this procedure.

After clicking the Run button, scroll through the list and select the logical cache type that you want to delete; then press Next to continue.

Add an Application Server

After clicking the Run button to perform this procedure, you will need to enter the following information when prompted:

Server Name

At this prompt, enter the name you want to use for the application server. This name is user defined. It can be whatever you want.

Is This a Windows Application Server?

If you are adding an Application Server to a system that has a UNIX root server, you will see this prompt: "Is this a Windows Application Server?" Answer Yes or No as appropriate for your system.

Note You will not see the prompt above if you are adding an Application Server to a system that has a Windows root server.

Network Address

The network address must be TCP. The TCP address is a four number. Each part of the address must be a number between 0 and 255 and separated by a period (for example, 135.0.20.91)

Delete Server

To delete a server, select the server you want to delete from the list box and click **OK**.

A message is displayed to warn you that you may be deleting a server that was configured as a system default server. Unless you cancel this operation, the system defaults for the server may be set to null. Since your system requires a default server, go to the System Defaults tab to reset defaults as required.

Add a Service to a Server

Services you can add include:

> Batch Service
> Cache Service
> SQL Service
> WorkFlo Queue Service
> Print Service
> ICR Service (ICR is NOT SUPPORTED in this Image Services release)

Note You cannot modify, or change, the print service that a printer is associated with using this procedure. Instead, you must first delete the printer and then add it again.

To add a service to a server, select these options from the scrolling lists:

- the server where the service will be added

- the service that you want to add

- dataset paths and dataset sizes

- use Fast Batch Committal option? (Y or N)

- the bes_commit process number

If the service you selected to add has already been configured on the server, a notification message displays. Choose another service to add.

Delete a Service from a Server

To delete a service from a server, scroll through the list of servers and select the server that has the service that you want to delete.

If only one server is configured on your system, the dialog box containing the list of servers does not appear.

Select the service that you want to delete and click Next.

If you select a service that is not configured on the server, a warning message appears. Click the Next button and select another service to delete.

Add a Storage Library Server

To add a storage library server to your system:

1 Enter the server name.

2 Enter the network address.

3 Enter the dataset path.

4 Enter the dataset size.

5 Select whether or not to use the fast-batch committal option.

6 Enter the number of bes_commit process.

Use the left and right arrows to enter a number for the bes_commit process. You can enter a value from 1 through 4.

Add VWService

The VWService you are adding must be associated with an SQL Service server number.

When you run this procedure, you will need to select a server to run the associated SQL service.

Modify VWService

This procedure is used to modify a VWService. Use this procedure to add WorkFlo Servers to, or delete them from, a VWService. The VWService you are modifying must be associated with an SQL Service server number.

When you run this procedure you will have to select the WorkFlo Management Service you want to modify, and the server(s) for that WorkFlo Management service.

Delete VWService

To delete a VWService, select the VWService from the list (if there are more than one) and click **Next**.

Index