

4.1.2



System Administrator's Handbook

4.1.2



System Administrator's Handbook

Note Sefore using this information and the product it supports, read the information in "Notices" on page 689.						
_						

This edition applies to version 4.1.2 of IBM FileNet Image Services (product number 5724-R95) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About This Manual 29

Document revision history 30

What to read first 30

Related documents 31

Conventions used in this manual 31

Flags 31

Multiple Server Configurations 32

Typing Instructions 32

Computer Output 32

File Paths 33

Cautions, Notes, and Tips 34

Command Syntax 34

Optional Parameters 34

Required Parameters 34

Accessing IBM FileNet documentation 35

IBM FileNet Education 35

Feedback 35

Documentation feedback 35

Product consumability feedback 36

Introduction 37

Image Services Servers 38

Combined Servers 38

Root/Index Servers 38

Storage Library Servers 39

Application Servers 39

Cache-Only Systems 40

Heterogeneous Systems 40

Peripheral Devices 41

Scan Servers 41

Print Servers 41

Fax Servers 42

Storage Libraries

Network Options 42

Networking 43

NCH and Three-Part Names 43

Object Name 44

Domain Name 44

Organization Name 44

Default Name 44

Software Services 45

Batch Entry Services 45

Cache Services 45

Index Services 46

Document Services 46

Print Services 46

Job Processing Priorities 47

Databases 47

MKF Databases 47

RDBMS Databases 48

```
RDBMS Naming Conventions 49
Transaction Logs 50
Cache Organization 50
Ageable Cache 51
Logical Cache Configuration 52
Cache Types 53
   Batch Cache 53
   Page Cache 53
   Print Cache 53
   Folder Notes Cache 54
Cache Backup 54
Entry Systems 57
FileNet P8 Content Federation Services for Image Services 58
Document Entry 59
Scanning 61
Verification 62
Document Assembly 63
Indexing 63
Committal 64
   Committal Modes 65
      Normal Committal 65
       Fast Batch Committal 67
   WorkFlo Distributor Queue Committal 68
   Transaction Logging 68
   Cross-System Committal 69
   Multiple-System Committal 69
Automating and Customizing Document Entry 70
```

Document Retrieval 70

Image Documents 71 Text Documents 71 Mixed Documents 71 Other Documents 72 Retrieval from Another System 72 New FileNet System Setup 73 System Serial Number 74 Software 74 Online Help 77 Directory Structures 77 Starting Image Services Software 79 Starting Application Executive for Image Services for UNIX Systems 80 Unified Logon for Image Services for Windows Server 81 Enable Unified Logon 82

2 **Database Maintenance**

Overview 85

Document Types

Media Families 86 Interleaving 88 Remote Committals 89 Library Servers 90 Write Surfaces 90 Indexes 91 Types of Indexes 92 Index Database 93

Disable Unified Logon 83 Customizing the Applications Menu 83 Index Database Tables 94

Retrieval Key and Informational Indexes 95

Cluster Indexes 97

Document Classes 98

Cataloging and Migration 100

Document Security 102

Auto Indexing 103

Expired Documents and Folders 106

Start Database Maintenance 108

Define a Media Family 109

Validate the Media Family 115

Add Disk Size 116

Change Disk Size 117

Commit to a Remote System 117

Define an Index 119

GUIDs Assigned to User Index 125

Add New GUIDS 127

Automatically Generate GUID 127

Manually Enter GUID 127

Rename GUID 128

Edit existing GUID 129

String Index 130

Numeric Index 132

Date Index 135

Menu Index 137

Create a Cluster Index 138

Assign a Cluster to a Document Class 141

Modify a Cluster Index 141

Delete a Cluster Index 141

Change Retrieval Key Status 142

Change a Retrieval Key Index to an Informational Index 142

Change an Informational Index to a Retrieval Key Index

Modify an Index 145

Rename an Index 147

Create a Menu Index 149

Change a Menu 152

Add, Modify, Delete Items in a Menu 152

Copy, Rename, Delete Menu 153

Create Document Classes 154

GUIDs Assigned to Document Class 167

Add New GUIDS 169

Automatically Generate GUID 169

Manually Enter GUID 170

Rename GUID 170

Edit existing GUID 171

Modify a Document Class 171

Change Indexes 172

Change Security 173

Save Changes 173

Update Retention Parameters 174

Select by Document Class 177

Select by Document Number 178

Select by Folders 180

Choose Retention Parameters 183

Start the Update 184

Update Document Security 185

Delete Expired Documents and Folders 188

Get Database Reports 190

Save Reports to a File 190

Print a Report 190

Find an Index, Document Class, or Media Family 190

User Indexes Report 191

Document Class Report 192

Media Family Report 193

Security Administration 194

Overview 194

Basic Concepts 195

Security Object 195

User 195

Group 195

Device 196

Permission 196

Administrator 196

Membership 197

Extended Membership 197

Administrative Domain 197

Session 198

Data Object 198

Override 198

Template 199

Expiration 199

Logging 199

Security Database 200

System Security 200 Group Security 201 Reserved Groups 202 Group Assignments 203 Administrative Group Assignments 203 Primary Group Assignments 204 Session Group Assignments 205 Group Deletions 205 User Security and Administrative Attributes 207 Classes of Users 207 Extensible User Authentication 209 Fallback Authentication 210 Document Security 211 Document Classes 211 Uncommitted Batches 212 Folders 212 FileNet Notes 212 Tabs 212 Function Security 213 Function Security Planning (RAC systems only) 213 Application Level Role versus Feature Level Role 217 Built-in Security Attributes 218 What Is a Role? 218 Relationships Between Function Codes and Roles 219 Default Behavior 220 Allow Access to Undefined Functions 220 Special user: SysAdmin 221 Best Practice 221 New System Considerations 221 Existing System Considerations 222

General Considerations 223

Case Examples 224

Device and Terminal Security 224

Password Controls 225

Logon Process 226

Start the Security Administration Application 227

Set Up FileNet System Security 228

Set Up System Defaults 228

Set Up the Default Group Template 235

Set Up the Default User Template 236

Set Up the Default Device Template 238

Override System Security Defaults 240

Override Logon Times 244

Override Security Object Defaults 245

Override Account Expiration Date 246

Reassign Objects to a Different Group 248

Plan Security Details 252

Example Scenario 253

Sample Setup 256

Set Up Group Security 259

Add or Update Groups 260

Delete Groups 265

To delete a Primary or Session group 265

To delete an Administrative group 266

Update Group Membership 266

Rename Groups 267

Using the Query Feature to Select a Group Name 268

Set Up User Security 269

Add Users 269

Set User Passwords 272

Extensible Password Authentication 272

Mandatory Password Change After Reset 273

User Expiration Exclusion 274

Custom Password Validation 274

Enable or Disable Custom Library Validation 275

Shared Library Entry Point 275

Extensible User Authentication 277

Enable or Disable Extensible User Authentication 277

Shared Library Entry Points 278

Fallback Authentication 280

Update Users 280

Changing the Name of the fnsw User 281

Delete Users 281

Update User Membership 282

Rename Users 283

Using the Query Feature to Select a User Name 284

Set User's Database Logon 285

Add Database Logon 286

Update Database Logon Password 286

Map Database Logon to Image Services Users 287

Delete Database Logon 289

Set Up Document Security 290

Set Up Function Security 290

Add Function Names 291

Activate Function Name 292

To activate a function name: 292

Update Function Membership 293

View Functions and Members 294

Deactivating a Function 294

Set Up Terminal and Device Security 295

Add Devices 296

Update Devices 297

Delete Devices 299

Update Device Membership 299

Rename Devices 302

Using the Query Feature to Select a Device Name 302

Task Flowchart 304

Security for Internetworking 305

Change Server Process Name 307 Change Server Process Password 308

Security Reports and Logs 309

Security Reports 309

Logon Reports 310

Event Logs 310

Locate a Directory or File 313

Save or Append the Event Logs 313

Print the Event Logs 314

System Management

System Management Functions 316

FileNet Task Manager 318

Monitor Functions 320

Status Information 320

Software Control Buttons 321

System Monitor 322

Monitoring the Image Services System 326

Monitoring Event Logs 327

Gathering FileNet System Information 330

Check the Active Processes 330 Monitor Network Activity 332

Monitoring the Flow of Work 334

Sessions 334 Batches 335 Print Jobs 336

Monitoring Storage Use 337

View Storage Use Information 338 Magnetic Disk Cache Information 338 Database Storage Information 340 Storage Library Information 341

View System Statistics 344

Networking Statistics 344

Document Services Statistics 345

Remove Unnecessary Media 348

Manage Other Files 348

Event Logs 348

Core Files in UNIX-Based Systems 349

Dr. Watson Files in Windows Server Systems 349

5 Database Server Connect 350

Overview 350

For DB2 only 351 For Oracle and SQL Server 352 For All Relational Databases 352

Database Connect Administration 353

Primary Area 353

Secondary Area (DB2 Only) 354

Changing the Primary and/or Secondary Password 355

Expiration Notification (DB2 only) 356

Password Maintenance 357

For Oracle and SQL Server 357

For DB2 358

Password Failure Emergency Procedures 359

For Oracle 359

For SQL Server 360

For DB2 361

Database Reconnect 362

Modifying the database reconnect default values 363

6 Background Job Control 364

Number of Background Jobs 366

Background Job Algorithms 366

Background Job Files 367

Checking Status of Jobs 367

Copying Media 368

Starting Background Job Control 370

Menu Bar 371

Current Jobs 372

Control Buttons 373

Importing Documents 374

Copy Annotations from Database to Media 376

Incorporate Media 378

Storage Libraries 379

For an optical storage library 379

For an MSAR library 380

Optical Disk Unit 381

Import Documents from Media 384

Verify Document Import 386

Copying Documents 389

Copy All Documents from a Surface to a Family 390

Copy Specific Documents from a Surface to a Family 393

Create a Document File List 393

Copy Documents Using a File List 395

Consolidating Media 398

Consolidate Media 399

Find Open Documents 402

Rebuilding Media 404

Erasing Media 406

Phases for Erase Media 407

For Optical Disk: 407

For MSAR Surface: 407

Erasing Media During Consolidation 408

To erase media during Consolidation 408

Erasing Media Without Consolidating 408

To erase media without Consolidating 408

Migrating Documents 409

Converting Optical Surface to MSAR Surface 412

Conversion Phases 413

Generating Reports 416

Completed Jobs 416

Results of Find Open Documents 418

Storage Library Control

Media 420

Preformatted 420

Surfaces 420

Transaction Log 421

Document Header Files 422

Storage Media Insertion (non-MSAR only) 423

Starting Storage Library Control 425

Messages 426

Informational Messages 427

RSVP Messages 428

Normal-Mode RSVP Messages 428

Manual-Mode RSVP Messages 429

RSVP/INFO Trigger 429

Message Display 430

Message Automation 431

Storage Libraries 431

Enable/Disable Library 431

Show Library Information 431

Backup Library 432

Library Management Functions 433

Show Library 433

Library 435

Pending Requests 435

Drives 436

Message Display 436

Enable/Disable Library 437

Calibrate Library 437

Enable/Disable Drive 438

Enable/Disable Slot 438

Enable/Disable Grippers 439

Media Management Functions 440

Preformat Media 440

Insert Media 441

Optical Storage Libraries 442

Optical Disk Units 443

MSAR Libraries 444

Eject Media 445

Storage Libraries 445

Optical Disk Units 446

Eject Media by Surface ID 446

Enable/Disable Media 447

Change Media Type 450

Change Media Family Name 451

Create Document Header File 453

Identify Media in Library 455

Reports 457

Slot/Drive Map 457

Local Statistics 458

Remote Committals 459

Pending Surface Requests 459

Detailed Surface Information 461

Specify Surface 461

Surface Description 462 Surface Statistics 462

Local/Foreign IDs 462

Media Surface Summary 464

Media Space Use 466

Media Family Information 467

8 Commands 468

Logging On to the Security System 468

Windows Server Users Logon 468 UNIX Users Logon 469

GUI Commands 471

initfnsw 471

vl 471

whatsup 472

Commands Overview 472

Arguments 473

Redirecting Output to a File 473

CSM exim 474

Determining What is Backed Up 475

Compressed Data 476

Running CSM_exim 476

Output Results 479

CSM tool 480

checkcache 481

listobjects 482

Statistics 484

Short Format 484 Sector Format 485 Long Format 485

cti 487

ddexim 488

deldocs 492

Document File List 493

Documents in Cache 493

Document Class Setting 493

Capture Committal Component 494

Erasable Media Surface 495

enlarge_ncol 496

Running the enlarge_ncol command 497

Example 498

fn_msg 500

gcp 501

ixdb_stat 503

less 509

PRI tool 510

cachestatus 511

cancel 513

checkcache 514

clearrequests 514

hardcopy 514

help 515

modify 516

printerstatus 518

requeststatus 521 resumeprinter 523 systemstatus 523 termoff 523 termon 524 spacerpt 525

Running Spacerpt on Servers with Remote Oracle Databases 529

ssn 529

stdocimp 530

Printing 534

How Printing Works 535

Print Services Database 536

Printer Selection 536

Request Ordering 538

Variable-length Pages 538

Print Services Initialization 539

Application and System Print Caches 540

Caching Strategy 541

Printing UNIX Files 542

Print File or Report Dialog Box 543

Interrupting a Print Job 545

Security 546

Troubleshooting 547

10 **Digital Document** Transfer System (DDTS) 552

Overview 552

Tape Drive Support 555

Maximums 556 Performance 556

Image Formats 557

CALS Files 558

Declaration File 559

Declaration File Records 560

Sample Declaration File 562

Raster Data Files 563

Export 565

The ddtsexp Command 565

Export Command Files 568

Global Options 569

Document Options 570

Page Options 572

Sample Import File 573

Import 574

The ddtsimp Command 575

Import Command Files 577

Global Options 578

Document Options 579

Page Option 581

Sample Import File 582

Creating a Script 583

Appendix A – Function Codes 586

Server Print Program

Database Maintenance 587

Storage Library Control 588

Background Job Control 589

Cache Export/Import Program 590

COLD Application

Overriding a Busy Batch

Full Functions Names Listing 592

Appendix B – Date and Time Formats 596

Current Date Mask Functionality 596

Support for Local Masks 597

Date and Time OS Settings 597

SUN Platform Date and Time Settings 598

Log File Date Format Exceptions 602

New Default Date Mask Application 603

Old Default Date Mask Application 604

Preserving the Old Date Mask on a Windows Server System 605

Preserving the Old Date Mask on a UNIX System 606

Conversion of 2-Digit to 4-Digit Date Format 607

Date and Time Number Conversion Ranges 608

Supported Date and Time Formats 609

Valid Separators 610

OS Exceptions 610

Masks without Separators 611

Default Formats 612

Windows Server Default Masks 612

UNIX Default Masks 613

Date Formats 614

Time Formats 615

Windows Server Date and Time Format Conversion 616

Appendix C – Logic for Retrieving Surfaces and Ejecting Media 619

Logic for Retrieving Surfaces 619

Logic for Ejecting Media 623

Appendix D – Task Manager Configuration File 624

Sample Configuration File 625

Appendix E – Message Triggering 626

Developing a Customized Script 626

Invoking the Script 627
Script Arguments 627
Error Tuples 631
Surface Ejection 632
Message Text 632
Media Types 635

Script Example 636

Appendix F – Multicultural Support 647

Getting started 647

Operating systems 647

Setting the Language variable 648 Setting the LANG variable on AIX 649 Setting the LANG variable on HP-UX 651 Setting the LANG variable on Solaris 652 Non-localized environment 653 Setting the LANG variable on Windows Server 653

Relational databases 654

Image Services configuration 655

Character set support 655 Single-byte character sets 655 Double-byte and multi-byte character sets 656 Heterogeneous environments 658 Known issues and limitations 659

Appendix G – Some Practical Limitations of **Image Services Components** 660

Overview 660

Document size 661 Storage Device Capacity 662 Network Speed Capacity 662

Glossary 663

Notices 689

Trademarks 692

U.S. Patents Disclosure 693

Index 694

About This Manual

Image Services System Administrator's Handbook describes the IBM® FileNet® Image Services system administrator's duties on the server and how to perform them. This handbook covers the full range of functions provided by Image Services Release 4.1.2. However, your system hardware and operating system (called a platform) or relational database management system (RDBMS) may not support all the functions and applications described here. For example, if you are using a combined server ssystem, none of the multiple-server or networking functions apply to your situation.

See also the *System Administrator's Companion for UNIX* or the *System Administrator's Companion for Windows Server.* To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Image Services is accessed using a desktop application such as *IDM Desktop, Capture*, and *Web Services/Open Client*. Throughout this manual, when we refer to your desktop application, we are talking about the PC client product you use to interface with Image Services.

If you have not purchased a particular FileNet software option, parts of this manual describing that option do not apply to your system.

Document revision history

IS version	Date	Comment
4.1.2	Nov. 2008	Initial release

What to read first

If you are new to the FileNet system, read Chapter 1, "Introduction," on page 37 in this manual, plus read the manual that came with your workstation software. For administrative functions on the workstation, see the PC coordinator's manual that came with your workstation software.

Chapter 1 provides an overview of the FileNet system components and checklists for implementing and administering the system.

The functions you use first in creating a new administrative setup begin in <u>Chapter 2</u>, "<u>Database Maintenance</u>," <u>on page 85</u>. This is followed by <u>Chapter 3</u>, "<u>Security Administration</u>," <u>on page 194</u> which describes how to set up your FileNet system security.

Chapters 4, 6, and 7 discuss ongoing administrative tasks and programs: System Management, Background Job Control, and Storage Library Control.

Chapter 8 provides command information including command syntax, command options, and using scripts.

Chapter 9 discusses printing from the FileNet system, and Chapter 10 explains the FileNet Digital Document Transfer System (DDTS) used to export and import documents.

Related documents

FileNet supplies the following system administration documents:

- System Administrator's Companion for UNIX
- System Administrator's Companion for Windows Server
- Enterprise Backup/Restore User's Guide
- MSAR Procedures and Guidelines
- FileNet P8 Content Federation Services for Image Services Guidelines

To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Conventions used in this manual

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

Flags

To help you identify information that applies to a specific Image Services server operating system, this manual uses the following labeled flags:



This flag identifies information that applies only to UNIX®-based Image Services systems.



This flag identifies information that applies only to Image Services for Windows® Server.

Multiple Server Configurations



This flag indicates information for users with more than one server. WorkGroup users and other users with single-server configurations should ignore sections with this flag.

Typing Instructions

To indicate commands, values, or other information you enter at your keyboard, we use the following indentation and typeface:

help [CSM_exim]

Computer Output

Information you see displayed at the console (such as displays of file contents, system messages, or output from program execution) is shown 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. *
```

Lengthy output listings are bounded by bold lines.


File Paths

This manual applies to all platforms. Where applicable, we include file path designations for both UNIX and Windows Server systems.

The term **FileNet local directory** refers to the directory for the specific type of system as follows:

/fnsw/local (UNIX)
\FNSW_LOC (Windows Server)

On UNIX systems, when you enter a file path in the File Selection dialog box, be sure to include a slash (/) at the end of the file path before you click the Filter button. A selectable list of files displays.



File Selection dialog box.

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 read.

Tip

Introduces an idea that might make your work easier.

Command Syntax

Command syntax definitions are indented:

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

Accessing IBM FileNet documentation

To access documentation for IBM FileNet products:

- 1 Navigate to the Information Management support page (www.ibm.com/software/data/support).
- **2** Select the appropriate IBM FileNet product from the "Select a category" list.
- 3 On the Product Support page, click **Documentation** and then click **Product Documentation**.
- 4 On the Product Documentation page, locate the document you need, then click the icon in the appropriate release column to access the document.

IBM FileNet Education

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

Feedback

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

Documentation feedback

Send comments on this publication or other IBM FileNet Image Services documentation by e-mail to comments@us.ibm.com. Be sure to include the name of the product, the version number of the product,

and the name and part number of the book (if applicable). If you are commenting on specific text, include the location of the text (for example, a help topic title, a chapter and section title, a table number, or a page number).

Product consumability feedback

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

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

1 Introduction

A FileNet system consists of client workstations that process work while connected to servers that run supporting services. Throughout this *System Administrator's Handbook*, **the system**, unless otherwise specified, refers to the set of software programs and hardware devices that make up the FileNet system.

The FileNet software that runs on the servers is called Image Services, see "Software" on page 74). On WorkGroup systems and small Visual WorkFlo systems, all Image Services software runs on a single server (combined). Larger systems might have the Image Services software distributed among multiple servers, with duplicate parts of the software on additional servers. Your system might not support all of the hardware and software functions described in this manual.

After your service representative installs and configures your system, you as the system administrator must keep the system running and notify your service representative when you need help (see <u>"System Management Functions" on page 316</u>).

Most FileNet system users run programs from PC workstations to enter, retrieve, print, and otherwise process documents. A few administrators use a special graphics-capable terminal, referred to as the **console**, to monitor and maintain the system.

Operators (nonadministrative users) log directly onto FileNet security using FileNet logon names. These users do not need operating system logon names to perform FileNet functions.

Administrators log into the native operating system of the server, start and stop the FileNet software, and log onto the FileNet software to run programs related to the FileNet application.

Administrators have an operating system logon name as well as a FileNet logon name, because FileNet security is separate from operating system security. Administrators create the FileNet logon names for all other users of the FileNet system, except for a few names and groups provided with the system. For details, see Chapter 3, "Security Administration," on page 194.

Image Services Servers

Image Services servers, the heart of the FileNet system, provide the central services that support document entry, document retrieval, and printing. (See <u>"Software Services" on page 45</u> for more information about IS services.) These servers also store the FileNet system's databases and provide cache storage areas.

Combined Servers

A single Image Services server that provides all services is called a Combined server. A Combined server is also referred to as a Root/Index/Library server.

Root/Index Servers



When you have two Image Services servers, you can separate the root and index functions from the storage library functions. The server that checks security, locates devices, and manages the index database is called the Root/Index server. For more details on databases, see "Databases" on page 47.

Storage Library Servers



In a system with two Image Services servers, the second server manages the storage libraries (optical disk jukeboxes or MSARs) and includes cache storage (on magnetic disk) as well as the related databases. Media is any material on which data is stored such as magnetic disk, optical disk and magnetic tape. When we refer to storage media, we are generally discussing optical disks.

The permanent database contains the storage media addresses of committed documents. The transient database tracks the cache contents, batch status, and other work in progress. Cache contains images on the way to and from storage media, other work in progress, and sometimes provides permanent data storage. For more information, see "MKF Databases" on page 47 and "Cache Organization" on page 50.

An Image Services system can have multiple storage library servers, each of which can manage up to eight libraries. If you have multiple storage library servers, one serves as the master server (called a **document locator server**) that keeps track of the contents of all storage libraries. Some diagnostic tools and commands must be run from the document locator server.

The document locator server includes the complete permanent database that maps each document number into one or two media locations. Other storage library servers contain small permanent databases that store this information until enough accumulates to transfer to the document locator server.

Application Servers



To extend processing functionality and media space, you can add Application servers to some systems. For example, Document Entry might need an extra server to process batches or manage a large number of Visual WorkFlo queues.

Cache-Only Systems

Cache is the magnetic disk space used to store documents on their way to and from storage media. In a system without a storage library or optical disk unit (a cache-only system), all documents reside in page cache.

As you add documents to your system, the index and permanent databases grow and consume more and more media space. You can only retrieve documents from page cache, so you need to regularly delete unneeded entries from the index and permanent databases to regain cache space. See "Delete Expired Documents and Folders" on page 188 for details.

See "Migrating Documents" on page 409 for details on migrating documents and "Change Disk Size" on page 117 for modifying media family disk sizes when you add a storage library to a cache-only system.

To help you decide which program to use for backing up cache, see "Cache Backup" on page 54.

Heterogeneous Systems

Image Services systems can be composed of servers that use a combination of different operating systems. However, the following limitations apply:

- An IS Root/Index server and Storage Library server must use the same version of the same operating system.
- An Application server that uses the Windows operating system can be configured with any UNIX or Windows IS system.

- A remote relational database server can use an operating system that is different from the IS Root/Index server and Storage Library server.
 - DB2 relational databases must always be installed on remote AIX or Solaris servers.
 - MS SQL relational databases can only be installed on local or remote Windows servers.
 - Oracle relational databases can be installed on either local or remote UNIX or Windows servers.

For example, within these limitations, the main IS system can be composed of UNIX servers, while a remote relational database server can be either a UNIX or a Windows server. Conversely, the main IS system can be composed of Windows servers, and a remote relational database server can be either a UNIX or a Windows server.

Peripheral Devices

The FileNet system supports many devices: scan servers, print servers, fax servers, and storage libraries. Other devices include tape, disk, and CD-ROM drives. All hardware devices depend on the Image Services server even if they also require additional servers.

Scan Servers

Scan servers are PC workstations equipped with scanners and scanning software. See the user's guide that came with the server for information about configuration, image formats, and operations.

Print Servers

All printers connect to a print server that receives documents from the Image Services. The print server also provides software to monitor net-

work connections, errors, statistics, and status information. See the user's guide that came with your print server for additional information.

Fax Servers

A fax server is a PC equipped with fax software and a fax card. The fax server receives inbound documents and converts them into FileNet images. It also forwards outbound documents to a fax machine or another FileNet fax server. See the user's guide that came with your fax server for more information.

Storage Libraries

A Storage Library server can control from one to eight Storage Libraries over a SCSI bus. A Storage Library is usually a jukebox for optical disks (storage media). Inside the jukebox, a robotic arm moves media between storage slots, the drives, and the operator's input/output slot as needed to satisfy requests.

A small system can use an optical disk unit (ODU). An ODU has no robotic arm; an operator manually inserts media into the optical disk drive and flips the media over as required.

Magnetic Storage and Retrieval (MSAR) media provides high speed and high capacity storage libraries on magnetic disk media.

Network Options

FileNet Image Services supports the TCP/IP (Transmission Control Protocol/Internet Protocol) network protocol. See the documentation that came with your desktop product for information about installing PC network hardware and software. An Image Services server can connect directly to any combination of four Ethernet, Token Ring, or Fiber Distributed Data Interface (FDDI) networks.

Networking



Two or more FileNet systems can be networked. In this way, workstations on one network can use the services and resources of other networked systems.

Administrative utilities such as Database Maintenance are available remotely through another console (using the X Windows protocol) on UNIX platforms. Administrative utilities are not available remotely on Windows Server platforms.

NCH and Three-Part Names

Because you can use services and equipment on any system networked to your own, you need a way to precisely identify what resource you want to use. The network clearinghouse (NCH) is a FileNet proprietary multi-keyed file (MKF) database that keeps track of resources and their addresses.

A resource is identified by a three-part name stored in the NCH database. The three parts of the resource name identify an object, a domain (system name), and an organization, in this format:

object:domain:organization

The maximum length of a three-part name is 82 characters—40 for the object, 20 for the domain, 20 for the organization, two for the colons separating the parts.

When you specify an object from a PC workstation, the maximum length of a three-part name is 79 characters—39 for the object, 19 for the domain, 19 for the organization, two for the colons.

Note

A DOS filename cannot exceed eight characters. If you transfer a file from a server to a PC running DOS, DOS truncates any name longer than eight characters when the PC receives the file.

Object Name

An **object** is a resource like a tape, printer, database, software service, logon name, etc. Some objects have names predefined by the system. For example, DefaultIMS is the name you use to access the index database and you cannot edit this name. Your service representative configures names for your printers and tape drives.

Domain Name

The **domain** is the system name. The system name is determined by you and set up by your service representative at FileNet system configuration time.

Organization Name

The third part of the NCH resource name is the **organization**. For example, you can enter your company name here.

Default Name



In most cases, the software displays a default name. To use a different resource (for example, service, device, or system), change the part of the name that is different.

For example, if a user on the "pubs" system wants to guery the main database on the "amber" system (within the same organization), the user need only change the middle name (the domain), as follows:

DefaultIMS:pubs:YourCorp

to

DefaultIMS:amber:YourCorp

Software Services

The FileNet software working behind the scenes is organized into services. The main services are batch entry services, cache services, index services, document services, and print services.

Batch Entry Services

Batch entry services manages the various phases of document entry: defining a batch, scanning, indexing, verifying, and committing; keeps batch information in the transient database; and works with cache services during batch entry and indexing.



In a multiple-server system, batch entry services normally runs on the storage library server. To reduce the workload on the storage library server in a heavily-used system, you can have additional copies of batch entry services running on application servers.

Cache Services

Cache services manages magnetic disk storage and retrieval of uncommitted objects, usually document pages.



In a multiple-server system, cache services can run on the storage library server and/or application server.

Index Services

Index services handles queries and updates to the index database. The index database stores document index records, document class and folder information, and the data dictionary.



In a multiple-server system, index services runs on the root/index server. The one exception to this is the case where the RDBMS is Oracle. In this case the index database can reside on a remote server.

In both cases another computer, even one not directly part of the FileNet system, can provide query, update, and index database *storage* functions.

Document Services

Document services (which runs on the document locator server in a mutiple-server system) manages document migration between cache and the storage library, document prefetching, and image ID allocation. Document services also initiates copy and import tasks. When given a document ID, document services locates the document, looking first in cache, then on media in a drive, and finally on media in a slot of the library. Only one copy of document services can exist on an Image Services system.

Print Services

Print services manages print requests for documents or text streams and displays information such as printer attributes, printer status, and job queue status. See Chapter 9, "Printing," on page 534 for more information about print services.



In a multiple-server system, print services usually reside on the storage library server. For improved printing performance at remote sites, you can include print services on application servers.

Job Processing Priorities

The FileNet system processes jobs based on priority. When you submit a print job, you can specify the priority of each job or accept the default priority.

See the table under <u>"Media Surface Summary" on page 464</u> for an explanation of storage media operation priorities.

Databases

The FileNet system includes several databases that are managed by either the FileNet multi-keyed file (MKF) subsystem or by a relational database management system (RDBMS) provided by a third-party vendor.

MKF Databases

The MKF subsystem manages the databases described in the following table. Some MKF databases can reside on different servers in a multiple-server system. If duplicated on another server, additional copies increment by 1 (represented by the **n** suffix in the database name).

Note

The MKF is limited to a maximum partition size of 16GB.

Database Name	Database Description			
Trans_DB <i>n</i> (transient database)	A directory of documents, images, and available cache space that tracks work in progress, including the status of batches, requests to read and write media, and print request queues.			
	Using information from the transient database, the system tracks images in cache to retrieve the images, when possible, from magnetic disk instead of storage media. The transient database tracks image use so it can delete the oldest images in cache to make room for new ones. Images not yet written to storage media are "locked" in the cache and cannot be deleted.			
Perm_DB <i>n</i> (permanent database)	Stores the media location of each document entered into the system. The permanent database also contains tables for media surfaces, media families, and notes (annotations, highlights and margin notes). The permanent database contains a row for every document committed to the system.			
NCH_DB0 (network clearinghouse database)	Serves as an address book for objects and services. Objects can be hardware (tapes and printers), databases, and programs (Image Services).			
Sec_DBn (security database)	Contains security information for each object (user, group, device), for the security service, for each deleted object, for each direct membership occurrence, and for each function name and class.			

RDBMS Databases

The index database and Visual WorkFlo queue databases are managed by an RDBMS. The index database is composed of a number of tables. The tables hold associated information, which is cross-referenced among the tables using common fields.

The four main tables of the index database are doctaba, user_index, document_class, and doc_class_index. Other tables are for folders,

cluster indexes, queues, and menus. The following chart describes the tables of the index database.

Table	Contents
doctaba	User index data for each document and FileNet system index information including the document ID.
user_index	Information about all user-defined index fields, including the index field name, data type, and whether the index is a retrieval key.
document_class	Information about every document class, including document class name, optional verification settings, and pages per batch.
doc_class_index	The indexes used by each document class.
folder	System-assigned folder numbers and user-defined folder names.
folder_contents	The document numbers associated with folders.
index_cluster	Location of cluster space on the system.
sys_numbers	Location of the next available document class number, folder number, and column name in doctaba.
WQM <number></number>	WorkFlo queues. <number> identifies the WorkFlo queue table.</number>
no_cat_audit	Log of when cataloging was turned on and off for each document class.
menu	A list of all menus associated with menu data types.
menu_items	A list of all menu items associated with their particular menus.

RDBMS Naming Conventions

FileNet databases managed by an RDBMS are discussed throughout this manual and the System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

You might be using one of several supported RDBMSs (for example, Oracle, DB2® or Microsoft® SQL Server™), so you will see references to the RDBMS in the generic format <yourRDBMS>. When you see this format, simply substitute the name of your RDBMS.

You can find a reference to the index database name in the following format:

<yourRDBMS_DB>

Transaction Logs

Databases maintain log files as sequential records of database changes. The log might be referred to as a redo log, a transaction log, a recovery log, or simply a log. Since these logs serve the same purpose, they can be generally referred to as transaction logs.

If the transaction logs reside on media other than the media that contains the database, the combination of restoring database backup tapes and transaction logs ensures a complete database recovery should one become necessary.


Your RDBMS maintains transaction logs for the index database. MKF maintains recovery logs for Sec_DB in Sec_RL1a, for Perm_DB in Perm_RL1a, and for Trans_DB in Trans_RL1a.

For detailed information about backing up databases and transaction logs, refer to your *System Administrator's Companion for UNIX* or *System Administrator's Companion for Windows Server* and your RDBMS manuals. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Cache Organization

Cache is the magnetic disk space used to store documents on their way to and from storage media. Within the physical disk space allocated to cache, cache is divided into **logical** caches. Using the System

Monitor's Magnetic Disk Cache Information report or CSM tool, you can list your system's logical caches with summary statistics for each. For more information about System Monitor report options, see "Monitoring Storage Use" on page 337. CSM tool is described in "CSM tool" on page 480.



Each type of data placed in cache is marked so that it is associated with a logical cache. For example, a scanned image is marked as an object in the batch cache (bes cache1 in the above sample list). When an operator commits the batch containing the image, the logical cache assignment for the image changes to the page cache (page cache1 in the sample). This remarking is referred to as a logical move because the data remains in the same physical location on disk.

Ageable Cache

Ageable cache is time-limited cache. Objects (usually retrieved documents) remain in this cache for a period of time specified during FileNet system configuration. After the time period expires, the objects are eligible for deletion if space is needed for another object. Locked objects (objects not migrated to storage media) do not age out. The system does not automatically delete locked objects, even though the time period has expired. To remove locked objects from cache, you must use the deldocs command. (For details, see "deldocs" on page 492.)

Logical Cache Configuration

Before your service representative configures your cache, you must estimate how many documents will be in each logical cache at one time.

Each logical cache has a minimum and maximum size (set during FileNet system configuration) which are percentages of the total available cache space. The minimum size guarantees that a logical cache has at least that much space. The maximum size is the upper boundary of the logical cache size. Logical cache sizes then contract and expand between these limits as needed.

Note

IS supports up to 255 16GB partitions, which allows for terabyte caches. The maximum cache size is 4080GB, or 4 terabytes. The maximum partition size depends on two operating system features: 1) The host operating system must support 16GB partitions; and 2) The host operating system must provide a mechanism which allows seeking to any offset up to 16GB from the beginning of a partition.

Cache Types

The FileNet system stores information in a number of different logical caches, each with a descriptive name that associates the cache with the type of information stored in it. Below is a brief description of each cache used in the FileNet system.

Batch Cache

This cache contains batches of documents as they enter the system, typically through scanning or from the COLD application. The batches remain in batch cache until they are ready to be committed. When the batch is ready to be committed, it moves logically from batch cache to page cache.

Page Cache

Page cache, also known as **retrieval** cache, contains all documents being committed to or retrieved from storage media. In addition, documents retrieved from media for printing are stored in page cache before being moved to print cache. Page cache is an ageable cache. Because an unlocked object remains in page cache until the space is needed, an ageable cache is usually full (within the limits of the configured minimum and maximum values).

Print Cache

Image Services uses two types of print cache:

- System print cache stores image documents waiting to be printed.
- Application print cache holds all other types of print job data, for example, text reports and files, waiting to be printed.

Folder Notes Cache

If your system includes FolderView, you can store folder notes in cache on the local PC, on a network PC server, or on a FileNet server.

Cache Backup

If you store items in cache that are not duplicated on storage media, you must back up cache. You can back up all of cache or just portions of cache.

Note

Cache Backup does not allow you to back up fast batch objects in page_cache. Cache Backup can see fast batch objects, but does not select them for export. If you are scheduling a back up of COLD objects only, the **Objects scheduled so far** value will read 0 because COLD always uses fast batch committal. All other locked objects in the cache are selectable and can be backed up. Cache Backup also exports all items in BES cache and imports them, but it skips all fast batch objects in the page cache.

Fast batch objects are always seen as objects in page_cache starting with ID number 4160000000. Cache Backup does not select any object with an ID greater than or equal to 4160000000. COLD objects will create temporary objects in page_cache with these numbers, plus any jobs that use the fast batch feature. To get around this issue, enable the **Fast Batch Breakup** feature on the System Application Services tab of fn_edit. After this feature is enabled, all new fast batch objects can be selected by Cache Backup, because Fast Batch Breakup makes these objects appear in the page_cache with the assigned doc_id rather than the temporary fast batch ID.

For more information about the fast batch feature, see "Committal" on page 64.

Items might be in cache and not on storage media for the following reasons:

- Your system does not have a storage library.
- You store folder notes in folder cache.
- You store Revise overlays in revise cache.

Before choosing a cache backup method, you should consider the following:

 Do you want to back up your cache to store for disaster recovery on the same Image Services server?

If so, you should use the FileNet Enterprise Backup and Restore (EBR) program. EBR synchronizes cache with the **transient** database, thus ensuring the restored objects have the same IDs they had when you performed cache backup. (For details on using EBR, see the *Enterprise Backup/Restore User's Guide*. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.)

 Do you want to back up locked objects to import on another Image Services server?

If so, you should use either the Cache Export/Import program or the CSM_exim tool. Both these programs use the same tape and file formats, allowing data interchange. However, because these programs do not synchronize the cache with the transient database, they are not the best tools to use for backing up cache to restore for disaster recovery on the same machine.

 Do you want to specify individual locked objects to export for import on another Image Services server?

If so, you should use the command-driven CSM_exim tool. If, however, you want to export a large number of objects (millions, for example), CSM_exim might fail with an out-of-memory error during export attempts since each object requires ten bytes of memory. For details, see the online help system or "CSM_exim" on page 474.

 Do you want to export all locked objects (such as documents) to another Image Services server?

If so, you should use the Cache Export/Import Program. This program provides an easy-to-use, graphical user interface (GUI) equivalent to the CSM_exim tool. However, unlike the CSM_exim tool, the Cache Export/Import Program does not export or import individually specified objects, read an ASCII file to specify objects, nor import cache data to a different cache.

See the online help system or the "Backup" chapter of your *System Administrator's Companion for UNIX* or *System Administrator's Companion for Windows Server* for general information about backing up cache and the Cache Export/Import Program. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Entry Systems



If you have to enter a very large volume of documents into a FileNet system, you can add one or more dedicated entry systems. Any server model can function as the combined server in an entry system. An entry system cannot have a printer and does not support WorkFlo.

The entry system must be compatible with the target system. Two systems are compatible when:

- Their document ID and surface ID numbers do not overlap.
- They share the same document class definitions.
- They share the same group definitions.

Your service representative sets up your Entry server so the document ID numbers and surface ID numbers do not overlap with the sequence of numbers assigned by the target system. If you have more than one entry system, the numbering sequence must not overlap that of another entry system.

In most cases, the entry system and target system use identical document classes. To ensure identical database definitions on both systems, perform initial database maintenance on one system (see Chapter 2, "Database Maintenance," on page 85), then import the database definitions to the other using the ddexim tool (see "ddexim" on page 488). If you change database definitions, make the same changes on the target system and all entry systems.

In addition to matching document class information, the two systems must have compatible security parameters. Committed documents acquire the access rights specified in the source system's document class and are committed to the media family specified on the source system, so you must be sure that identical group names and media family names exist on the target system. The users in those groups can differ, but the read access group name, the write access group name, and the append/execute access group name must be present on both systems (see Chapter 3, "Security Administration," on page 194).

The remote entry system (RES) is supported for Image Services and Visual WorkFlo applications. A RES is a single server configured to commit documents to a remote system over a high-speed communications link. A Remote Entry Server has no storage media of its own. A RES and target system must have compatible document ID numbers. Surface ID numbers are irrelevant because a RES has no surface ID numbers of its own. A RES supports document entry through scanning.

FileNet P8 Content Federation Services for Image Services

FileNet P8 Content Federation Services combines the capabilities of Image Services (IS) systems and Content Engine (CE) systems, allowing content stored on the Image Services system to be cataloged and viewed on a Content Engine System.

With FileNet P8 Content Federation Services, new documents, images, and other content are entered either from the IS system or the CE system and stored in specific document classes within Image Services. Existing Image Services index properties in existing document classes can also be exported to the Content Engine system.

FileNet P8 Content Federation Services provides Record Management functionality. You can delete a document as well as declare and delete a record.

There are several ways that a FileNet P8 Content Federation Services system can be configured.

- Documents entered into a document class on the Content Engine system are only cataloged on the Content Engine system, so even though the documents are stored in Image Services, they can be retrieved only by the Content Engine.
- Documents entered into a document class on the Image Services system can be indexed on both the IS and CE systems, enabling documents to be retrieved from either system. Or you can choose not to keep the index properties in Image Services, so only the Content Engine can retrieve the documents.
- Existing documents in a document class on the Image Services system can have their index values mapped to corresponding properties in a document class on the Content Engine. You can choose to delete the indexes from Image Services so the documents can only be retrieved by the Content Engine.

To configure a FileNet P8 Content Federation Services system, contact your service representative about purchasing the additional FileNet software needed to enable CFS on the Content Engine system.

Document Entry

Document entry is the process by which documents are entered into the system.

Note

You can now use FileNet client applications to import a wide variety of image and non-image data, such as audio, video, e-mail, and Electronic Data Interchange (EDI) documents.

Entering documents using FileNet client applications involves a number of steps, some required and some optional. The number of optional steps, as well as the order in which they are performed vary, based on your particular installation.

Document entry includes the following required steps:

- Capturing the paper or electronic pages by scanning pages or importing files.
- Assembling electronic images of pages into electronic documents.
- Committing batches of electronic documents to the Image Services system.

Document entry might also include any of the following optional steps:

- Verifying that scanned images are of acceptable quality.
- Processing the document to improve image quality.
- Indexing document properties and attributes of electronic documents.
- Verifying that index values of document properties and attributes are correct.

In the paragraphs below, we describe each document entry step, as well as the services and general concepts associated with each step. In this discussion, we present document entry through scanning from a document entry station (a workstation with an attached scanner).

Before you can set up batches for scanning, the system administrator creates a document class for each type of document using the Database Maintenance application. The document entry operator must select the appropriate document class for each batch of documents.

Batches have read, write, and append security permissions you set up when defining a document class. Any user can set up the batch. However, all users who scan documents into the batch, index them, or commit them must have full access (read, write, and append/execute permissions) to the document class.

For Image Services/WorkGroup systems, the scan station provides document entry functions. For information about the scan station, see the appropriate client documentation.



In a multiple-server system, when you use a document entry station on one system (system A) to commit documents to another system (system B), the document entry station operates like a remote station on system B. When the operator selects the batch service on system A, the software logs the operator into system B. The operator must use a logon name that is valid on both systems. See "Extensible Password Authentication" on page 272 for information about setting up user logon names for networked systems.

Scanning

The document entry process assigns a number to each scanned image. The first image number becomes the document ID number unless you delete the image from the batch or rearrange the pages.

Each document has a header that contains its document number, document class, indexing data, pointers to the rest of the pages in the document, and other information. The permanent database contains

the media address of this header, which is stored as page 0. When you look at the list of objects in page cache, you will notice that the size of page 0 is much smaller than other pages with the same document number (object id).

<csm_tool></csm_tool>	>1 ssn	object id	page	max_length		
1	10291	20026080	0	232	ageable	
1	10291	20026080	1	359889	ageable	
1	10291	20026080	2	359889	ageable	
1	10291	20026080	3	359889	ageable	
1	10291	20026080	4	359889	ageable	
1	10291	20026080	5	359889	ageable	
1	10291	20026080	6	359889	ageable	
1	10291	20026080	7	359889	ageable	
1	10291	20026080	8	359889	ageable	
1	10291	20026080	9	359889	ageable	
1	10291	20037321	0	116		
1	10291	20037321	1	6562		
1	10291	20037323	0	124		
1	10291	20037323	1	3064		
1	10291	20037462	1	3072		
'CR' = lir	ne, 'space	' = page, '	p' =]	paging off,	'x' / 'q' = exit	

Verification

Image verification is an optional step between scanning and document assembly, in which the scanning operator verifies the scanned image on the screen. The operator can rescan any unacceptable images. Other verification options (such as index verification and batch total verification) are described in the "Database Maintenance" chapter (see "Verification Options" on page 162).

Document Assembly

Document assembly organizes the individual scanned images contained in the batch into documents, which can then be indexed and committed to page cache. Depending on the type of documents in the batch, as well as the structure and requirements of the repository, you can assemble documents using different techniques.

Operators can assemble documents as they are scanning them, or offline, after completing the capture process. If a very fast scanner is available, you can maximize throughput by having assembly functions performed at another workstation. This would optimize performance of the scanner workstation.

Indexing

When you define a document class, you select the indexes associated with that class. These indexes specify the information that will later be used to retrieve the documents. For example, suppose you create a document class that uses an index for social security number. The screen form displays the name you assigned to the index, prompting the indexing operator to enter the value—the social security number that appears on the scanned document.

Note

In a FileNet P8 Content Federation Services environment, a document class might or might not have index values associated with it. For more information about defining a document class, see "Document Classes" on page 98.

In many cases, indexing operators key in index values for each scanned document. Operators can also scan bar codes that are automatically translated into index values. (See <u>"Automating and Customizing Document Entry" on page 70</u> for more information.)

The FileNet software automatically stores some index information, like document ID, date of entry, document class, and media family name. You decide what additional information to store for each type of document. When different kinds of documents require different types of information (for example, different indexes or a different number of pages) or they have different security requirements, you create different document classes to accommodate each.

All index information is stored in the index database and also on storage media in page 0 of the document. When you retrieve the document, the FileNet software searches the index database for index information that satisfies the retrieval query.

Committal

Committal is the stage at which scanned images become documents. Committing a batch of images includes the following:

- Assembly: The creation of a document in page cache.
- Cataloging: writing the index values to the index database (doctaba table).
- Migration: storing the document on storage media and writing the storage media addresses to the permanent database (docs table).

When the document images are written to storage media, transaction logging might also occur (see "Transaction Logging" on page 68 and "Media Families" on page 86). However, depending on your environment, some of these tasks can be omitted. For example, you might choose to commit your documents without migrating them to storage media, as in the case of systems that do not have storage libraries.

Committal Modes

The following is an overview of normal and fast batch committal modes and general committal concepts.

Normal Committal

Normal committal takes place when scanning and indexing are complete for a batch of documents. Committing the documents in a batch includes some or all of the following steps:

- Create document IDs.
- **2** Catalog documents (optional in WorkFlo systems).
- **3** Move documents from batch cache to page cache (making the documents available for retrieval).
- 4 Migrate documents from page cache to storage media (optional).

Creating documents is an internal process in which cache services moves images from batch cache to page cache, and updates permanent database tables.

Cataloging is the process of writing information to the index database. The documents are available for retrieval from page cache once the indexes are in the index database. You can turn off cataloging if you store indexes elsewhere than on the Image Services server. Some sites prefer to store index values in a different database, usually on another computer system. Cataloging can be set system-wide during configuration of your FileNet system. For more information about cataloging, see "Cataloging and Migration" on page 100.

You can retrieve documents as soon as they are catalogued, regardless of their migration status. Documents residing in page cache can be retrieved very fast. When you request a document, the IS software

looks for that document in page cache first. If the document is not found in page cache, a check for the document on storage media follows.

Migrating is the process in which the FileNet software copies documents from page cache to storage media and writes storage media addresses to the permanent database.

You can choose several migration options when you set up document classes (see "Create Document Classes" on page 154):

- migrate the documents as soon as possible
- delay document migration until a predetermined time
- never migrate the documents

At committal, workstation users with proper security can override the established document class migration options.

Your service representative can configure the FileNet software to give priority to read or write operations. By default, the FileNet software gives priority to read requests so retrieval operators get optimum response to document retrieval requests. Other lower priority tasks include writing documents to storage media.

To improve the performance of media writes, the software uses a scheduling scheme that minimizes the number of robotic arm movements and media insertions. When storing documents according to media family rather than cluster index, the software performs all writes to one surface before it starts writing to the next surface.

Fast Batch Committal



In a multiple-server environment, fast batch committal is a quick way of moving a large number of images to storage media. In a fast batch committal environment, all documents and all pages of the documents for the batch are in one cache object, so the entire batch is committed in a single operation. COLD always uses fast batch committal.

Note

With fast batch committal, either the entire batch is migrated to storage media or the entire batch is not migrated. This process ignores the migrate/no migrate option you set in the document class (see "Cataloging and Migration" on page 100).

Your service representative can configure your FileNet system to use fast batch committal exclusively. However, you should not use fast batch committal for migration when operators need to access documents immediately after committal, because the index information is written only **after** writing all storage media copies.

The drives doing mostly fast batch committal should favor writes over reads.

Note

If fast batch committal is configured, you cannot use cluster indexes. See "Cluster Indexes" on page 97 for details on clustering.

WorkFlo Distributor Queue Committal

When you create a document class in Database Maintenance, you can associate it with a WorkFlo distributor queue.

For each document committal in this document class, the system adds document IDs to the queue, even if cataloging is disabled. After committal, a WorkFlo program usually examines the queue entries for these documents and routes them to the appropriate user (based on document ID or user index values) for processing.

When configuring your FileNet system, your service representative determines whether or not to include user index values in the WorkFlo distributor queue entries.

Transaction Logging

When you use transaction logging, the system makes at least two copies of a document, writing each copy to different storage media. One copy goes to the primary media. One or more additional copies go to transaction logging media. The system writes to all transaction logging media before writing to the primary media.

You can use additional transaction logs to export documents to a remote system or as backup media for disaster recovery. To save room in the storage library and to provide extra security, most sites store filled transaction logging media at another location.



In a multiple-server system, **remote committal** is an alternative to physically moving multiple transaction logs to another location. Using remote committal, you commit to a primary media family on a remote (networked) system. See "Multiple-System Committal" on page 69.

Cross-System Committal



When defining a batch, you select a batch entry service (BES) that controls scanning, indexing, and committal. You can select a batch service on your own local system or on a remote system.

Normally, a batch entry service (BES) commits documents to the local system. Alternatively, your service representative can configure a batch service to commit the documents to another system (for example, a remote Image Services system). This is **cross-system committal**. For more information about configuring these systems, see the *Multi-Committal and Cross-Committal Configuration Handbook*. To download this handbook from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Multiple-System Committal



If you need copies of documents on more than one system, you can link a primary family on the local system to one or more primary families on other systems. Remote committals start after all local committals are finished and documents remain in cache until all remote committals are finished.

The difference between cross-system committal and multiple-system committal is in the number of primary family copies of the committed documents. With cross-system committal, documents are only committed to one primary family that resides on a remote system. With multiple-system committal, documents reside on the local system as well as on one or more remote systems. Any of these systems can use transaction logging to have two or more copies of the documents—one on the primary family's media and one on each transaction logging media. For more information about configuring these systems, see the

Multi-Committal and Cross-Committal Configuration Handbook. To download this handbook from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Automating and Customizing Document Entry

In an Image Services system, you can automate and customize your document entry processes using the capabilities of FileNet Capture Professional and Workflow products.

Document Retrieval

The Image Services system can act solely as a document repository. However, you usually store documents because you want to use them later, perhaps retrieving a document so you can respond to a customer's question. When a paper document is stored in electronic form, anyone who has the proper security can display a copy of the document without waiting for others to finish with it. If you need a paper copy, you can print or fax it.

- For more information about document retrieval, see the documentation for your specific FileNet client application.
- For more information about the logic involved in retrieving a document, see "Appendix C – Logic for Retrieving Surfaces and Ejecting Media" on page 619.

Document Types

Regardless of the method used, you can retrieve any supported document type. Image Services supports several document types, including image, text, form, mixed, and other (non-displayable).

Image Documents

Image documents can be one of the following:

- A document scanned, indexed, and committed with WorkFlo/Scan
- A document received from a fax machine and processed by a fax server
- A document created through the UNIX Visual WorkFlo/Scan Window call, which saves a snapshot of a foreign window as an image
- A Windows object (a document committed by a non-FileNet Windows application, such as Word for Windows)

Text Documents

Text documents include:

- COLD documents having no background images
- FileNet/WorkFlo/Fax journal logs

Mixed Documents

A mixed document can be:

- A COLD document with a background image
- A tiled image (a large image that is divided into rectangular sections called tiles)
- A document entered through the IXF/Sun Import application
- A WordFlo document

A document containing more than one page type

Other Documents

Other is a document type reserved for documents that cannot be displayed or printed. This includes archived DOS files stored on the FileNet system storage media.

Retrieval from Another System



If your FileNet system is networked to another FileNet system, you can retrieve documents from both systems (with proper security).

For example, from a PC, you can log on to system A, then retrieve images from system B. You must have the same name and password on systems A and B, and have read access to the document. When you retrieve an image from system B, the system moves the image to system B's retrieval cache, then to the workstation.

When printing documents from remote systems, the system transfers documents across the network and into a print cache on the local system. This prevents a print service on one system from filling up a cache on another system.

New FileNet System Setup

As System Administrator for a new system, one of your first tasks is to set up the FileNet system.

Chapter 2, "Database Maintenance," on page 85 and Chapter 3, "Security Administration," on page 194 provide details on setting up and maintaining the system.

Chapter 4, "System Management," on page 315 and Chapter 7, "Storage Library Control," on page 419 provide additional information about the tools and procedures for ongoing administration of the FileNet system.

The following checklist provides an overview of suggested procedures that are required to implement and maintain a FileNet system.

Security	Database Definition	Ongoing Procedures	Application Software
☐ Create security matrix	☐ Define media families	☐ Manage security administration	☐ Enter documents ☐ Conduct queries
☐ Determine system defaults ☐ Define administra-	☐ Define cluster indexes* ☐ Define index fields	☐ Perform backups ☐ Monitor media space	☐ Create and maintain forms
tors, users, groups Set up object security	☐ Define document classes	☐ Monitor storage media activity☐ Monitor printing	☐ Implement FolderView* ☐ Create WorkFlo
☐ Set up function security	Define archive Image Services servers*		scripts*
☐ Secure system accounts	*optional		*optional

System Serial Number

Each Image Services system (domain) requires a System Serial Number (SSN) that not only identifies the specific IS system, but also helps identify all the records stored within the system.

This identifier is vital to Image Services, as the SSN is included in the metadata of each committed record. Each IS system must have its own unique SSN to configure peer systems in a cross-system committal environment, and to prevent potential problems if media are ever transferred from one IS system to another.

When installing a new Image Services systems, you are required to construct your own SSN using any combination of numeric digits that you choose. The valid range of SSNs is 1000 through 2147483646.

Note

Current Image Services systems will continue to use their existing SSNs, which will always be valid.

Software

The Image Services software consists of applications and services. Most applications appear on the Application Executive's Applications menu. Applications forward user requests to FileNet software services.

For example, when a user issues a print request, the application passes the request to print services. If your system is connected to another FileNet system, you can select services on that system as well as your own. See "Software Services" on page 45.

The primary applications needed for system administration are:

System Administration Applications

Application	Description
Task Manager	Starts and stops FileNet software and displays event logs. See <u>"FileNet Task Manager" on page 318</u> for information about this program.
	On UNIX platforms, booting the system automatically starts the FileNet software.
	On Windows Server, you must use the Task Manager to start FileNet software, then log onto FileNet security and get access to the applications described below.
Database	Sets up and maintains your index database.
Maintenance	See <u>Chapter 2, "Database Maintenance," on page 85</u> to create and modify indexes, document classes, and media families and to archive Image Services servers (a collection of archive databases you can query).
	You can use the reporting features of Database Maintenance to obtain information about user indexes, document classes, and media families. Database Maintenance also includes several special-purpose tools.
Security Administration	Creates user accounts for the FileNet software and sets up other security attributes (see Chapter 3 , " Security Administration ," on page 194).
System Monitor	Displays information and statistics about Image Services resources. Image Services system, security, and services information displays automatically.
	You can select options to display specific information about storage components (magnetic disk, cache, and storage libraries) and network and document services statistics. For more information about the System Monitor, see "System Monitor" on page 322.
Background Job Control	Starts functions (copying media, importing documents from another system, consolidating media, etc.) that normally run in the background.
	You can get information about both current and completed jobs, as well as suspend, restart, and terminate current jobs. See Chapter 6 , "Background Job Control," on page 364 for details.

System Administration Applications, Continued

Application	Description
Storage Library Control	Monitors media requests and displays other messages related to your storage libraries.
	In a multiple-server system, you run Storage Library Control on the consoles attached to the storage library servers. See "Starting Storage Library Control" on page 425.
Cache Export/Import	Backs up your cache storage independently of your system backups.
	You should use this program to back up locked objects such as documents only if you want to export your cache data to import on another Image Services server. Because this program does not synchronize the cache with the transient database, it isn't the best tool to use for backing up cache to restore for disaster recovery on the same machine.
	See the System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server for information about the Cache Export/Import Program. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Note

On a Remote Admin Console (RAC) system or a server whose logon domain is remote, the following menus will be disabled (grayed out): Archive Utilities, Update Document Security, and Update Retention Parameters.

For more information, see the Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Online Help

All GUI interfaces have online help. The help system provides a table of contents, keyword search capabilities, hypertext links, and browse sequences.


The Help pulldown 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 for 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

Directory Structures

The FileNet software resides in a distinct file system—fnsw—to avoid conflicts with other standard file systems.

The following illustration shows the fnsw subdirectories in a typical directory structure. Depending on your particular system, your directory structure might look slightly different.




This file system exists, in similar but not identical form, on every UNIXbased Image Services server.

For the Image Services for Windows Server platform, refer to your System Administrator's Companion for Windows Server for information about the directory structure. To download this document from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Starting Image Services Software

Before you use any Image Services functions, the FileNet software must be running. Use the Task Manager to determine the state of the FileNet software.



If FileNet software is running, a long list of processes displays in the middle of the window. Become familiar with these processes; missing processes can indicate a problem.

You can stop or restart the software by clicking the appropriate button. If FileNet software is not running, you see only the TM_daemon process. In this state you can start the software or enter backup mode.

Important

Do not run IS commands or start IS processes while shutting down the IS software. Starting IS processes while IS is trying to shut down could cause the system to hang. If you create automated scripts or cron jobs that run IS processes, ensure that they are disabled before you shutdown IS for maintenence or backup. Verify that monitor programs are not scheduled to start during backup times.

Starting Application Executive for Image Services for UNIX Systems

After the FileNet software has started, start the Application Executive by double-clicking the icon in the FileNet Image Services Server Application program group or by entering **Xapex** at the command line. If you have system administrator privileges, log on as SysAdmin (password: SysAdmin). Otherwise, log on as Operator (password: Operator).

After you answer the logon prompts, the Application Executive main window displays.






After you successfully log on to the Application Executive, it checks your logon credentials to determine your user access to applications. When the system has validated your logon, you can choose a program from the Applications menu.

To exit the Application Executive, click the Logoff button on the main window, then select Exit from the File menu.

Unified Logon for Image Services for Windows Server



After the FileNet software has started, you can start the Application Executive by selecting it from the Start menu path Programs/FileNet Image Services Server Applications. This dialog appears:



Note

This Application Executive window is different for Remote Admin Console installations. See the *Remote Admin Console User's Guide*. for more information. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Enable Unified Logon

To set up the association between your Windows Server logon and your FileNet logon:

1 Enter your Image Services user logon and password.

If you have system administrator privileges, log on as SysAdmin (password: SysAdmin). Otherwise, log on as Operator (password: Operator).

- 2 Check the **Update cached credentials for automation** checkbox.
- 3 Click OK.

The next time you log onto Windows Server, you are also automatically logged onto Image Services and do not need to open Application Executive.

If you log onto Windows Server under another user name, you need to open Application Executive and enter this information for the new name.

You can choose an Image Services program from the Application Executive's Applications menu or from the Start menu path Programs/FileNet Image Services Server Applications.

Note

When you make changes to a user's password in the Security Administration program, you must then run the Application Executive and reset that user's password information for unified logon.

Disable Unified Logon

To delete the unified logon association:

- 1 Click the Application Executive's Logoff button.
- 2 Check the **Delete cached credentials for automatic logon** checkbox.
- 3 Click OK.

The Windows Server logon name is no longer associated with the FileNet logon name.

Customizing the Applications Menu

You can customize the Application Executive's Applications menu by creating the file .apexcust.cfg in the appropriate format. The location of the file is based on the system user's (not the Image Services user) HOME environment variable.

Note

Windows Server users can create and set the HOME environment variable in the Control Panel in System Properties/Environment/User Variables.

If it doesn't find the file in the home directory, the system searches the directory /fnsw/bin (UNIX) or the directory specified by the FNSW LOCAL DRIVE configuration item (Windows Server). If found, the file is parsed based on the following format:

cprogram_name>:cprogram_style>:[cprogram_title>]

where:

program_name> identifies the program to run.

style> is the string gui or tty, identifying the user interface
 style. On most platforms this string is required to start the pro gram.

am_title> identifies the string to display on the Applications
 menu. This field is optional. If not defined, program_name>

A # character in the first column represents a comment — the system ignores the rest of the line.

Note

Do not provide the full path name of the file. The system expects to find the program in /fnsw/bin (UNIX) or FNSW_LOCAL_DRIVE (Windows Server). If you provide the full path name, the entry appears grayed out in the Applications menu.

Database Maintenance

Overview

You use the Database Maintenance application to:

- Create and maintain media families, indexes, and document classes
- Change indexes from primary to informational, or from informational to primary
- Delete documents and folders
- Update document security
- View reports on indexes, document classes, and media families

One of the first administration tasks on a new system is to create the media families, indexes, and document classes (in that order) needed to organize documents.

For Remote Admin Console (RAC) users, you can select Database maintenance options by clicking the function button or by using the menus. The options tied to the function buttons are:

- Define/Update Index
- Define/Update Family
- Define/Update Class

- Indexes
- Families
- Classes

For more information about the Remote Admin Console, see the *Remote Admin Console User's Guide*. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Media Families

The media family defines what type of storage media (optical disks or MSAR) the document class uses.

In general, the media family controls which media surfaces the document classes use. You may also use the media family to configure one or more remote systems for multiple system committal.

Note

Even if you do not use storage media (you keep everything in cache on magnetic disk), you must create a primary media family before you can create a document class.

The system needs access to media both for writing (committal, copying, etc.) and reading (document retrieval, printing, etc.). The more media families you use, the more drives you should have. For example, if you have two drives and often use three media families, the storage library would constantly swap storage media to insert the correct media for the family into the drive. Excessive media swapping not only slows performance, it also prematurely wears out the hardware.

A **primary family** is a collection of media that stores documents from one or more document classes. Before you can complete the definition of a document class, you must assign it to a media family so the system knows where to store documents in that class.

The FileNet system is configured to write one or more extra copies of committed documents to **transaction log** media. You can define one or more media families for transaction logging. The order of images on transaction log media is the order in which images are committed, regardless of the primary media family name and media type.

You can define up to eight transaction log families, giving you that number of backup copies of a committed document. You can use these backups to distribute extra copies of documents to other sites or store media off site for disaster recovery. With multiple transaction logging, you can create a new tranlog family for the new media type if, for example, you move from 2.6 GB media to 7 GB media.

You can enable or disable transaction logging for each primary family. At committal time, the local primary media and all its transaction log media must reside in storage libraries attached to the same library server. For remote committals, the remote systems must be available. For retrievals, primary and transaction log media can reside in any storage library. The system writes only the addresses of the primary family and one transaction log (the principal transaction log) in the permanent database. When you assign transaction logs, the last entry assigned is the principal transaction log.

Note

Documents from a secondary transaction log surface must be imported into the system (including the system on which it originated) before you can use them. Once imported, the secondary transaction

log surface becomes the principal transaction log for those documents on its surface (but not for the primary family in general).

For added protection of transaction log copies, you may use a media rotation scheme for a secondary transaction log. For example, if you fill individual transaction log media every five days, you might institute a cycle that fills 10 media (5 principal media and 5 secondary media) in 25 days. You start principal transaction log media at the beginning of each week for five weeks. You change the secondary transaction log media daily, using one for each day of the week.

The pattern looks like this:

Rotation Week	Monday	Tuesday	Wednesday	Thursday	Friday
Week 1	P1 and S1	P1 and S2	P1 and S3	P1 and S4	P1 and S5
Week 2	P2 and S1	P2 and S2	P2 and S3	P2 and S4	P2 and S5
Week 3	P3 and S1	P3 and S2	P3 and S3	P3 and S4	P3 and S5
Week 4	P4 and S1	P4 and S2	P4 and S3	P4 and S4	P4 and S5
Week 5	P5 and S1	P5 and S2	P5 and S3	P5 and S4	P5 and S5
P=Principal Transaction Log Media S=Secondary Transaction Log Media					

Interleaving

Interleaving is the process of writing the A sides of several storage media before returning to write the B sides. Interleaving media keeps the most recently committed documents immediately available for retrieval without flipping the disk to write to the other side. For more information about interleaving, see the steps under "Define a Media Family" on page 109. You need at least one more drive than the number of interleave surfaces to make interleaving useful. One drive is for the transaction log media and the other drives keep the surfaces

available for committal and retrieval. Do not use interleaving when importing documents from full media to another system (filling media takes longer when interleaving is enabled).

Remote Committals



To commit documents to another system in addition to the local system (see "Commit to a Remote System" on page 117), you link a primary family on the local system to one or more primary families on other systems (the destination system could link to a third system). You cannot use remote committal to write to an archive database.

Note

Do not define remote committals if you use delayed migration or no migration. Remote committals do not happen until after the local committal, so page cache could fill quickly if you use remote committal in conjunction with delayed migration. In the case of no migration, the documents would never commit to the remote system.

When you commit a batch to the local system, the system commits the first family in the list. When finished with that, the system commits the next one. If it doesn't find any family, the system starts committal over at the beginning of the list, which could cause duplicate committals. Likewise, if you reboot the system, committal starts over at the top of the list. You can change the family commit list at any time.

Remote committal uses fast batch committal (see <u>"Fast Batch Committal" on page 67</u>). If the remote system is compatible, the document numbers remain the same on the two systems. If the remote system is not compatible, it must generate new document IDs, slowing down the committal process.

Library Servers



At committal (on the local system), the primary media and all its transaction log media can reside in different storage libraries, but the libraries must all be attached to the same library server. You can assign a media family to a particular library server and storage library (see Step 6 on page 110).

Write Surfaces

You set how many storage media surfaces in a family can be available for writes at one time (see steps under "Define a Media Family" on page 109). The default number of surfaces is one. By specifying more than one current write surface, you can balance the load of writes to storage media during normal committals and also during media copies through Background Job Control (see "Copying Documents" on page 389). Specifying two current write surfaces requires three drives (just for writing) when using transaction logging.

The system attempts to balance the load of writes by writing first to one surface and then to another. If you need the drives for retrievals, however, the system does not use two write surfaces until the retrieval demand subsides. If your system has a heavy document entry load, multiple current write surfaces can help speed up the process. Multiple current write surfaces also make better use of multiple storage libraries. The default number of surfaces is one on each eligible server. If you assign a primary family to specific servers or storage libraries, you can assign the current surfaces to some or all eligible storage libraries.

If you use only one transaction log (tranlog) and make any assignments to the tranlog family, you must assign at least one current write surface for every library server on the system. This is necessary so the

system can write committed documents to a transaction log on the same server on which the primary family resides.

Indexes

Indexes are labels associated with a document (for example, social security number, account number, name, address). Index values allow you to retrieve documents stored in cache or on storage media.

Note

In a FileNet P8 Content Federation Services environment, indexes are not required when documents are captured on the Content Engine. If the documents are captured on the IS system, indexes are required even though they may not be stored on the IS system.

You choose the kinds of information you want to use when retrieving documents and create indexes to gather that information. Index operators then enter a value for each index during document entry.

You can create up to 224 indexes. You cannot delete an index, but you can rename an index to reuse it. Each index you create increases the size of the database and the amount of time needed to back up the database.

An index name typically consists of 1-18 alphanumeric characters in any case (upper, lower, or mixed), without spaces or any other special characters, such as the Euro € character, except for an underscore.

An index should **not** have:

A name starting with the characters F (capital F, underscore). These names are reserved for FileNet system use.

- A name that starts with a, b, d, f, p, pv, q, t, w, x, or X followed by a number (for example, b2, f591, q7648, t99, or w573). These can all be names of system-generated document entry batches.
- A name that starts with wfl followed by an underscore and numerals (for example, wfl 44734).
- Any of the following reserved words: div, mod, like, and, not, or, find, next, prior, via, in, range, having, more, of, where, keywords, defined, image, text, form, mixed, other.
- A name that is an Image Services or Visual WorkFlo reserved word. For more information, see your client documentation.

Types of Indexes

The following table describes the four types of indexes.

Index Type	Description
Numeric	Use the numeric type for index values used in calculations. You can run batch totals only on a numeric field.
	A numeric index value can, when input, have a maximum of 30 characters. On output, it can also contain commas, a plus or minus sign, and a dollar sign, based on the mask you provide.
String	Use a string index to store identifying numbers or number and letter sequences (such as social security or account numbers) or any index entries that do not fit into any other category.
	A string includes 1 to 239 characters. You cannot run batch totals on a string even if all of the characters in the string are numerals.

Index Type	Description
Date	Use a date index to cause the system to check index entries for valid input.
Menu	Use a menu when so few entries exist that you can list them on the indexing form and let the indexing operator choose the value instead of keying it or selecting it from another dialog box.

Index Database

FileNet uses a relational database management system (RDBMS) to store much of your data. The primary RDBMS database is called the index database (see also "RDBMS Databases" on page 48). The index database stores:

- Document index values, index field definitions, and document class definitions
- DMA properties, including a display name and from one to ten GUIDs for each user index. (For details, see "GUIDs Assigned to User Index" on page 125.)
- An audit trail showing each time index cataloging is turned on or off for all document classes
- Menus
- Folder definitions, and folder contents, and FolderView records
- Visual Workflo queues

On a Visual Workflo queue server, an RDBMS database stores additional Visual Workflo queues.

Index Database Tables

The four main tables of the index database are doctaba, user index. document class, and doc class index. Other tables are used for folders, cluster indexes, queues, and menus. The RDBMS manages the index database. See the chart under "RDBMS Databases" on page 48 for a description of the tables of the index database.

Every document in doctaba has one entry (row) in the table. Each row in doctaba represents one document and doctaba associates a document's index values with the document ID. Other tables use the document ID as a cross-reference. In doctaba, 31 columns are reserved for FileNet index fields, and 224 columns are for user-defined index fields (see "Define an Index" on page 119).

The doctaba table includes the following FileNet fields for which the system automatically generates values.

FileNet Field	Description
F_DOCNUMBER	System-assigned document number.
F_DOCCLASSNUMBER	System-assigned document class number.
F_ENTRYDATE	Date the document was cataloged.
F_ARCHIVEDATE	Date the document can be archived.
F_DELETEDATE	Date the document can be deleted.
F_DOCTYPE	Document type: image, text, form, or mixed.
F_DOCFORMAT	Document format: text type, image type, etc.
F_DOCLOCATION	Document location: whether the document is located internally or outside the FileNet system.


Retrieval Key and Informational Indexes

You must indicate whether each index you create is a retrieval key. You may use only a retrieval key as a primary key in a query. A retrieval key should produce only a small number of matches. Good candidates for retrieval keys are social security number (SSN), loan number, and account number.

A retrieval key has an extra structure called a B-tree. A B-tree is a table containing sorted index data that matches each retrieval key entry (such as a social security number) to its row ID number in the data records table. Then, when you search for a retrieval key, the system finds the specific element in the B-tree, then finds the correct row in the data records.

Any index that is not a retrieval key is considered an informational index. You can use an informational index to retrieve documents, but since the data is stored in no particular order, the system has to look at each document sequentially to see if it has the requested information. This takes significantly longer than a search based on a retrieval key.

The following illustration shows the difference between the structures of a retrieval key index and an informational index. A retrieval key index search starts in the smaller B-tree structure, then proceeds to the correct row in the data records. An informational index search starts at the first row and goes down until it finds the correct information.



You may combine secondary (filter) conditions with a retrieval key. For example, a query might request all documents for persons named Brown (a retrieval key) born in 1946 and living in Chicago (two informational fields).

If you retrieve documents using a primary key in combination with one or more secondary filters, the system first collects document numbers that qualify according to the primary key (which must be a retrieval key). The system then looks sequentially in this reduced set of documents to find those that qualify according to the secondary filters.

Cluster Indexes

Clustering directs the system to store all documents with a common index value (for example, the same loan number) in a reserved space on particular storage media. To do this, you tell the system how many documents you are going to store for that index value, and the system saves space on media for future documents.

Note

Only string and numeric retrieval key indexes can be cluster indexes.

The primary reason to cluster documents is the need to access all or most of the documents associated with one index value at the same time. For example, you might scan a set of documents over a period of time, filling several storage media. If you store the documents on any media you are using at the time, then need to access all of the documents at the same time, the system would have to read documents from different media. This causes delays in retrieving all the documents you need because the system might have to move several storage media in and out of a drive to find all the documents.

If you think your application might benefit from clustering, discuss it thoroughly with your service representative. If you reserve more space than you need, you'll waste space on storage media. If you reserve too little space, the documents end up on different storage media in spite of clustering.

To use clustering effectively, the following must be true:

- You can accurately predict the average number of documents in a cluster.
- Documents are consistent in size.
- The number of available storage media does not exceed the capacity of the storage library. (If you need to access 500 different media but only have one 200-media library, clustering won't be useful.)
- You often need to retrieve an entire cluster of documents with no prior warning, so overnight prefetching of documents is not a viable alternative.

Note

Clustering cannot be used if your system is configured to use fast batch committal (see <u>"Fast Batch Committal" on page 67</u>).

Document Classes

All documents that enter the FileNet system, regardless of the entry method (scanning, tape files, disk files), must belong to a document class. When you set up a document class in Database Maintenance, you determine certain defaults for all documents in that class, such as number of pages per document, media family name, whether to verify images, who can access the documents, and which indexes to use.

Note

The maximum number of document classes that can be defined in an Image Services system is 256.

Define separate document classes when your requirements for scanning and indexing differ for different types of documents. For example,

you would use different document classes when you use different indexes in each. Also define different document classes for documents that have different security restrictions.

A document class name typically consists of 1-18 alphanumeric characters in any case (upper, lower, or mixed), without spaces or any other special characters, such as the Euro € character, except for an underscore. You cannot start a document class name with a numeral or with the characters F_ (capital F, underscore).

In addition, a document class name should **not** have:

- A name that starts with a, b, d, f, p, pv, q, t, w, x, or X followed by a number (for example, b2, f591, g7648, t99, or w573). These can all be names of system generated document entry batches.
- A name that starts with wfl followed by an underscore and numerals (for example, wfl 44734).
- Any of the following reserved words: div, mod, like, and, not, or, find, next, prior, via, in, range, having, more, of, where, keywords, defined, image, text, form, mixed, other.
- A name that is a client software reserved word. For more information, see your client software documentation.

When storing documents on storage media, you can organize them in two ways:

- Fill up storage media with documents all from the same document class, in which case you just assign the document class to a media family.
- Put all related documents together at one place on storage media. You might scan the documents under different document classes,

but need to keep them together at one place on storage media. This is called clustering. You can read more about clustering in "Cluster Indexes" on page 97.

Cataloging and Migration

Cataloging is the process of writing index information to the index database during committal. Sites that keep index information on another computer can disable cataloging on the FileNet system. Your service representative can configure your system to disable cataloging system-wide or let you disable it for certain document classes. If you disable cataloging, you can still add documents to a Visual Workflo queue.

The term **migration** means writing to storage media. When you have storage libraries, you can choose to turn off migration for one or more document classes (see also "Fast Batch Committal" on page 67).

You have three migration options:

- Migrate immediately after committal
- Migrate at a specified time (delayed migration)
- Do not migrate

Note

With fast batch committal, either the entire batch is migrated to storage media or the entire batch is not migrated. This process ignores the migrate/no migrate option you set in the document class (see "Fast Batch Committal" on page 67).

Important!

In a multi-Storage Library server environment, if you do not choose to migrate documents from cache immediately after committal, (that is

you choose Migrate at a specified time or Do not migrate) make sure that a preferred Storage Library has been assigned for the Media Family. Otherwise, when you decide to migrate the documents, you may receive an error message. To resolve, assign a preferred library for the family.

When you do not write documents to storage media, they remain in page cache. Documents require a lot of media space. You should configure your system with enough cache space to contain all the documents you want to keep on magnetic disk.

Documents are vulnerable to loss when they exist only in cache. If you do not write documents to storage media or you delay migration, you should back up the entire cache using the tool most appropriate for your system requirements:

- To back up cache for restoring on the same Image Services server in case of disaster, use the Enterprise Backup and Restore program. For details, see the Enterprise Backup/Restore User's Guide.
- To export cache to import on a different Image Services server, use the command-driven CSM exim tool or its GUI-interface equivalent, Cache Import/Export Program. See "CSM_exim" on page 474 and your System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server.

To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

When you set delayed migration, the system catalogs the documents and deletes the batch, but delays the final steps of actually writing the documents to storage media and updating the permanent database with the address. When the time you set for delay expires, the system

automatically writes the associated documents that remain in cache to storage media and updates the permanent database.

You cannot force documents to storage media earlier than the time specified. If you change the delay for a document class, that change does not affect documents already written to cache with a different delay time.

Document Security

The document class controls document security. The groups assigned to a document class determine whether all users, an individual user, or a group of users can read, write, or append/execute documents assigned to the document class.

The default security setting is (ANYONE). If you do not change access, any user may scan, index, commit, and retrieve documents in a document class (subject to functional limitations or other security options).

The program used to export database definitions (ddexim) sets the security for all document classes to (ANYONE)/(ANYONE). Be sure to set security appropriately on the stand-alone entry system before you start scanning.

The following table summarizes what operators can do with various access rights to the document. The "X" indicates the type of privilege required for an operator to perform the corresponding action.

Operator Action	Read	Write	Append/ Execute
Retrieve	X		
Display and print document	Х		


Operator Action	Read	Write	Append/ Execute
Create annotations	X		Х
Delete documents	Х	Х	Х
Change attributes	Х	Х	
Scan, index, and commit	X	Х	X
Reassemble pages	Х	Х	Х
Delete uncommitted batches	Х	Х	Х

Any restrictions you specify apply also to uncommitted batches. Any operator can set up the batch, but only those with all three access rights can scan, index, and commit batches that use this class.

Auto Indexing

For string indexes only, you can choose to set up automatic indexing for your document class.

When you select auto indexing, the following information appears in the lower section of the Add/Update User Indexes dialog box (see Step a on page 165).



Auto Indexing Data

In addition to document images, a scan station can capture a stream of text that you can direct to indexes. You specify how to read index values from this text stream in your auto-indexing definitions. This eliminates the need for some or all manual indexing.

The person who sets up the scanner defines text segments that supply the information you require in the order that you specify here. Available data includes bar code data, scanner data (such as image ID and operator name), and server data (such as system time). Although bar codes are normally used to separate documents, they can supply a stream of text that you can use for indexing.

You must work with the person who is setting up the scan station to coordinate the information you require with what the operator can supply. Since the scan station must work under the control of a document class, you must provide the setup requirements first. In general, you can consider one text segment delivered by the scan station to be the equivalent of one bar code.

You can use up to five lines to specify the location and size of an index. For example, assume that the scanner delivers four text segments as shown below:

```
1234567890123456789012345... (Column Counters)
  1412 E ELM STREET
  APT 14B
  COSTA MESA, CA
  92626
6740599 222551111
```

From these text segments, you can read values for the indexes named Address, Account_Num, and Soc_Sec_Num. In the column titled Barcode #, enter the text segment number, as defined in your client software documentation.

In this example, the value of the Address field comes from four different text segments. The software concatenates four pieces of text into one index value. Be sure to include any necessary spaces. The index value for Address will be:

1412 E ELM ST APT 14B COSTA MESA CA 92626

Expired Documents and Folders

As you add documents to your system, the index and permanent databases grow and consume more and more disk space. You can delete the entries for unneeded documents and folders from the index and permanent databases to regain magnetic disk space. (If the documents are stored permanently in cache rather than on storage media, deleting removes committed documents from cache.) To regain even more disk space, you can delete folders and unfile documents from folders after a specified period of time.

Prior to deleting expired documents and folders, you should check the way document classes are set up. Make sure the retention base (Date Filed or Date Closed) and time (Months from) fields are set correctly and disposition is Delete. Note that merely changing retention parameters for a document class does not change documents already committed to that class. To change those documents, you must use the Update Retention Parameters option.

An alternative to deleting records from the databases is to archive them. Archiving copies the records to an archive database on storage media before deleting them. If you do need to retrieve an older document, you can query the archive database instead of the magnetic disk (default Image Services) database.

Note

On a Remote Admin Console (RAC) system whose logon in remote, archiving is disabled. For more information about RAC, see the Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

While the document images remain on storage media, deleting their indexing information (their means of being located) removes the capability to access the documents. After deleting most of the documents from storage media, you can use Background Job Control's Consolidate Media option (see "Consolidate Media" on page 399) to copy the remaining active documents to other storage media and remove the older media from the library. You can erase media after consolidating if you use erasable media.

The function for deleting expired documents and folders deletes records in the databases for folders and documents flagged as expired (if you choose both folders and documents). An item expires when the specified number of months have passed after the retention base (date filed or closed).

The process first unfiles any documents from folders if the document unfiling date has expired, then deletes obsolete folders and obsolete documents. If a folder expires before the documents in it, the process unfiles the documents and deletes the folder. It does not delete the documents until they too are obsolete.

Start Database Maintenance

Choose Database Maintenance from the Application Executive's Applications menu. The Database Maintenance main window appears.

Menu options are described in the following table (see also "Online" Help" on page 77). Use this table as a quick reference when you need to find a particular Database Maintenance function.

Database Maintenance menus

Miscellaneous	Indexes	Classes	Families
Archiving Utilities	Define/Update Index	Define/Update Class	Define/Update
Delete Doc./Folder	Rename	Report	Family
Update Doc. Security	Build Retrieval Key		Report
Update Retention Parameters	Drop Retrieval Key		
Exit	Define/Update Cluster		
	Report		

CAUTION

It is highly recommended that the Define/Update options under both the Indexes and Classes menus be executed during non-production hours. These tools will issue locks on the relational database (Oracle or SQL Server) that may cause long delays in system processing. This primarily affects systems with very large databases, but is the recommended procedure for all systems.

The Report function under both the Indexes and Classes menus provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.


CAUTION

Any modifications made to the Index fields, Document Classes and Disk Families will not take effect until the FileNet software is recycled.

Define a Media Family

You must define a media family before you can create document classes, even if you keep all your documents in magnetic disk cache.

Choose Define/Update Family from the main window's Families menu 1 to display the following window.



- 2 Click the List button beside the Family Name field to see a list of existing family names.
- 3 Select a name for updating, or enter a new name and click OK.

To enter a new name, type only alphanumeric characters and underscores (up to 18 characters) in the blank field beneath the list.

A media family name consists of 1-18 alphanumeric characters in any case (upper, lower, or mixed), without spaces or any other special characters, such as the Euro € character, except for an underscore.

4 Select the disk size (if you are creating a new family).

For a system without storage media, you can select any disk size. If you later add storage media, you can change the disk size to match the media (see "Change Disk Size" on page 117).

Click the arrow to display the list of media types. For any media family name (whether primary or transaction log), you must choose one type of media. However, the transaction log types can be different from the primary type.

Note

If you enter a size that is not configured, an error message does not appear until you attempt to save the family.

5 Select the interleave count.

> To set up interleaving, click the appropriate button to the left of the number. You can change this number after creating the family.

6 Click the Yes radio button in the Preferred Library field to assign more than one current write surface to speed up committals or to specify writing to a particular server or library. The system will then only look at the storage library before performing a write.

This displays the Currently Assigned Surfaces list box along with the Assign and Delete buttons.

Note

Always ensure that the preferred library setting for the family is still correct after adding or deleting storage libraries.

Important!

In a multi-Storage Library server environment, if you do not choose to migrate documents from cache immediately after committal (see "Cataloging and Migration" on page 100), you must assign a preferred storage library for the media family before migrating documents.

Important!

When the Image Services system is put into production, the system administrator should monitor cache resources frequently to prevent the cache on a particular server from becoming full. If the cache on a Storage Library server nears capacity, the system administrator can reassign the Preferred Library of a media family to prevent it from filling up completely.

If you make no assignments, the system attempts to balance the writing load. The system assigns one current write surface to the next eligible library server (a server with the proper media type for the family).

For example, if you have a multiple storage library configuration where the primary family and an associated tranlog is write compatible on one server but not on the others, then the writes will go to the servers with the primary/tranlog setup. However, if your multiple Storage Library server configuration is such that the primary family and an associated tranlog are write compatible on all servers, the system will automatically balance the writes among all servers. Please note that Image Services requires that the primary and its associated tranlog surfaces reside on the same Storage Library server. Also, be aware that the system will look at all storage libraries before determining which one(s) to write to. Be aware of the following rules:


- If a primary family does not have a preferred library and it has an associated tranlog family, then the primary family and tranlog family must be write compatible on the same Storage Library server.
- If a primary family has a preferred library and it has an associated tranlog family, then all tranlogs must be write compatible on the same preferred server.

For a system without storage media, this selection is ignored.

Note

If you have a multiple storage library server configuration, click the Yes option in the Preferred Library area and assign a preferred storage library. If the storage library has an associated tranlog family then it must be attached to the same storage library server.

- 7 Click the Assign button to display a list of library servers.
- 8 Select a library server and click the Assign button to display the following dialog box (in the illustration, one write surface is already selected).



Note

The purpose of this tool is to divide the writes up among assigned surfaces so you can optimize the writes between the surfaces. You can divide writes among different surfaces on the same storage library or among different surfaces on multiple storage libraries, depending upon the surfaces you have assigned. If you have several surfaces assigned, the tool will write to the first available surface to speed up the access time.

- Click the down arrow to display a list (1–8) and select the number of write surfaces you want to be active at one time.
- Click the appropriate button in the grid to assign a surface to a particular storage library. For example, to assign surface 1 to library A, click the upper left button (1 A). To assign surface 2 to library B, click the second button on the second line (2 B).

c Click the OK button when you are finished, then close the Storage Library Server dialog box by clicking Cancel. The list box on the main window displays your assignments.


Note

If you assign a surface to a non-existent library, the system displays this message when you try to save the family:

Cannot assign current surface to a Storage Library which does not exist or does not match the family media type.

- Select the family type. 9
 - If you are defining a primary media family, click the Primary radio button and skip to step 12.
 - If you are defining a transaction log media family, select the TranLog radio button and go to step 10.
- 10 Assign a transaction log family to the primary family.
 - To assign a transaction log family to a primary family:
 - a Click the Add button, which displays a list of transaction log families.
 - b Select one or more names from the list and click OK. The names appear in the list on the main window.
 - To write to several transaction log media, choose one family whose addresses you want to retain in your permanent database and be sure that name is the last one on the list.
 - To add a tranlog family later:
 - temporarily delete the name of the tranlog family from Step a

- add the new tranlog family
- add back the tranlog family you deleted
- 11 When you finish all definitions, choose Save from the File menu.
- 12 Confirm your choice of media at the prompt.



13 Click OK to dismiss the popup window that confirms you created the family.

Validate the Media Family

Whenever a change in Storage Library configuration occurs, especially when a Storage Library is deleted, you must validate the media families by resaving the media families manually and resolving any errors.

Follow these steps to resave a media family:

- 1 Choose Define/Update Family from the main window's Families menu.
- 2 Click the List button beside the Family Name field to see a list of existing family names.
- 3 Select the name of the media family you wish to save and click OK.
- 4 Choose Save from the File menu.

If the system saves the media families successfully, you're done.

If the system does not save the media families, you will see an error message indicating the library number that is incorrect. Some reasons for save errors include:

- deleting a library
- reassigning a library number
- configuring a new library type over (with the same letter as the old one.)
- adding a new (OSAR) server and more.

Correct the warnings and errors that appear by adding or changing the preferred library to match the current storage library configuration.

Add Disk Size

In an existing system with a storage library, you can add a new drive that supports a different disk size. When you do this, you must create new media families for this new disk size.

CAUTION

Do not modify MKF database tables to add a new disk size.

To add a new disk size to an existing system:

- 1 Commit all outstanding batches.
- 2 Create new media families for both primary and secondary surfaces.
- 3 Add the new families to a document class.

All future batches for this document class go to the new disk type.

Change Disk Size

After you create a media family, you can change the disk size if the surfaces are not assigned. This is usually done if you add a storage library to a cache-only system (all committed documents are stored on magnetic disk) so you can migrate documents to the storage library.

CAUTION

Do not modify MKF database tables to change the disk size.

To add storage media to a cache-only system:

- 1 Commit all outstanding batches.
- 2 Change the disk size for your media families for both primary and secondary surfaces.
- 3 Use Background Job Control to migrate already committed documents from magnetic disk to storage media (see "Migrating Documents" on page 409).

Commit to a Remote System



Before you can commit to a remote system, your system must be configured correctly (with a remote domain) and that system must be available. You must also be logged on with a user name that is valid on both your local system and the remote system.

- Choose Define/Update Media Families from the Database 1 Maintenance window's Families menu.
- 2 Select the primary family whose documents you want to commit to a remote family.

- 3 Choose Families from the Remote menu to display the Remote Media Families dialog box.
- 4 Click the down arrow beside the Domain text field and select the system to which you want to commit the documents.

Note

For Remote Admin Console (RAC) users, when you select the target domain, you will see a pop-up window prompting you to enter an account password for the user common to both systems. For more information about RAC, see the Remote Admin Console User's Guide.

Click the down arrow beside the Families text field and select the family 5 on the remote system that will receive the documents.

Note

For Remote Admin Console (RAC) users, when you save the family you just added, you will see a pop-up window prompting you to enter an account password for the user common to both systems. For more information about RAC, see the Remote Admin Console User's Guide. To download the RAC User's Guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

6 Each family you choose on the remote system appears at the top of the screen when you click the Add button. When you finish assigning families at one site, you can select another site and add families.

Note

If the remote family uses transaction logging and the remote site does not want to make additional copies of the documents, the remote site must delete the transaction log families from that primary family.

When you are finished with all assignments, click OK. 7

Define an Index

Indexes must exist before you can create document classes.

Note

In a FileNet P8 Content Federation Services environment, a document class may or may not have index values associated with it. If you want the documents associated with the document class to be retrievable only from the Content Engine system, indexes are not required. However, if you want the documents to be retrievable from both the Content Engine system and the Image Services system, you need to specify indexes in the Document Class. For information about mapping the index values between the IS and the CE systems, see the FileNet P8 Content Federation Services for Image Services Guidelines. To download these guidelines from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

CAUTION

Do not define unnecessary indexes. The maximum number of indexes you can define is 224. You cannot delete an index or change its type. For restrictions on changing an index, see "Modify an Index" on page 145.

Before you can define a new index, you must make sure that no one else accesses the database. While other users are on the system, you can modify and get information about existing indexes without disturbing normal system operation.


However, you may lose any data you enter for a new index. If any index database activity (committal, retrieval, filing documents in folders, etc.) occurs before you save the new index, the system cannot write the data and an error message appears.

Tip

To prevent users from logging on while you're defining new indexes, put all the users in a session group and expire the session.


To define an index, follow these steps:

1 Choose Define/Update Index from the Database Maintenance window's Indexes menu, which displays the following window.



The Define User Indexes window displays a warning about the time required to create a retrieval key. To change an index to a retrieval key later, see "Change Retrieval Key Status" on page 142.

- 2 Define the index name.
 - Click the List button to enter a new name. This displays the User Indexes dialog box listing all current indexes.




To view information about an index, select its name and click the Ok button. Check the list to see if the index name you want to use already exists.

- Enter a new name in the Selection field at the bottom of the dialog box. This name appears on the default indexing forms of document classes that use the index. See "Indexes" on page 91 for naming conventions.
- Click the Ok button to save the index name. The list box closes and the cursor moves to the description field of the Define User Indexes window.

- d Enter a description for the index that clarifies its purpose. You see the description in this window and in the User Indexes Report. It does not appear in list boxes. A description can have from 1 to 30 characters, including spaces and special characters, in any case (upper, lower, or mixed).
- Define the DMA properties. The DMA Properties enable you to specify 3 a DMA display name and up to 10 GUIDs (Globally Unique IDentifiers) for each defined user index.
 - Specify the Display Name. When you first enter a name in the Index Name field, the system automatically duplicates that name in the DMA Properties Display Name field. You can change this name by selecting all or part of the text and typing the desired name.
 - This field accepts up to 30 alphanumeric characters, including spaces and special characters.

b Click the Edit button to open the GUIDS dialog box, which enables you to add, delete, or modify GUIDs for the currently defined user index. You must assign at least one GUID to the index and may assign up to 10 GUIDs for each user index.



- Select the GUID you want associated with the user index. For details on adding, editing, and deleting GUIDs, see "GUIDs Assigned to User Index" on page 125.
- d Click Ok. The selected GUID displays in the GUIDS field on the Define User Indexes dialog box. Each additional GUID defined on the GUID list may serve as an alias for the user index.

- Define the index type. 4
 - Choose an index type. An index is one of four types: numeric, string, date, or menu. Additional parameters appear at the bottom of the window as appropriate for the index type.
 - Indicate if the index is a retrieval key by selecting either the Yes or No radio button.
 - Select Yes if you want to use the index as a primary key (instead of only using it as a secondary filter) in a guery. You may use only a retrieval key as a primary key in a query.
 - Select No if you never want to use the index as a primary key. Secondary filters always use a sequential search.
 - c Complete additional parameters for the type of index. See "String Index" on page 130, "Date Index" on page 135, "Numeric Index" on page 132, or "Menu Index" on page 137.

Note

Once you have created an index, you cannot delete it or change its type.

5 Save the index.


> When you complete all fields for each index, choose Save from the File menu. If no other database activity or input errors occur, the system saves the index successfully. Otherwise, it displays an error message. If this happens, you'll need to try later or change the appropriate fields.

GUIDs Assigned to User Index

GUIDs (Globally Unique IDentifiers) are DMA-compliant, 16-byte integers used to uniquely identify each element transported over a network. The system ensures unique GUID assignments by automatically generating this integer using an algorithm based on the system's network card MAC address and a format that complies with the specifications provided for the system's platform.

Each GUID must have a unique name that conforms to the format specified for your Image Services system platform. For platform specifications, see the GUID naming conventions described in the documentation that came with your operating system.

To add, delete, or rename GUIDs for the current user index, click the Edit button in the DMA Properties box of the Define User Indexes dialog box. The GUIDs dialog box opens, displaying a list of currently defined GUIDs.



As you add or rename GUIDs, the system validates the assigned name, making sure it meets the following requirements:


- Unique name assignment
- No duplication of name in the Image Services database
- At least one and no more than 10 GUIDs assigned to each user index

You may modify the currently defined list as follows:

- To delete an existing GUID, select it from the list and click the Delete button. A message box opens, asking you to verify deletion. Click Yes to delete the selected GUID or No to cancel deletion.
- To add a new GUID to this user index, click the Add button. The Add New GUIDS dialog box opens. For details, see "Add New GUIDS" on page 127.
- To rename an existing GUID, select it from the list and click the Edit button. The Edit GUID dialog box opens, displaying the exiting name of the selected GUID. For details, see "Rename GUID" on page 128.

Add New GUIDS

The Add New GUIDS dialog box opens after you click the Add button on the GUIDS dialog box when defining the DMA properties of a user index.



Automatically Generate GUID

Click this radio button to have the system automatically generate a GUID. The system displays the GUID in grayed-out, uneditable text.

To save the automatically generated GUID, click the Ok button. The system closes the Add New GUIDS dialog box and displays the newly assigned GUID in the GUIDS list.

Manually Enter GUID


Click this radio button if you want to manually assign a GUID. Type a 16-byte character set into the field, using the format shown in the sample displayed below the entry field.

To save the manually defined GUID, click the Ok button.

- If you entered a GUID that already exists in the Image Services database or does not use the correct format, a message box opens, alerting you to the problem. When this happens, you must enter an acceptable GUID and click Ok again.
- If you entered a unique GUID using the correct format, the system closes the Add New GUIDS dialog box and displays the newly defined GUID on the GUIDS list.

Rename GUID

The Edit GUID dialog box opens after you select an existing GUID from the GUIDS list and click the Edit button on the GUIDS dialog box when defining the DMA properties of a user index.



Edit existing GUID


This field displays the currently selected GUID and allows you to rename it.

To change the current name, select all or any part of the displayed text, type in the desired text, and click Ok.

- If you entered a GUID that already exists in the Image Services database or does not use the correct format, a message box opens, alerting you to the problem. When this happens, you must enter an acceptable GUID and click Ok again.
- If you entered a unique GUID using the correct format, the system closes the Edit GUID dialog box and displays the renamed GUID on the GUIDS list.

String Index

Selecting the String index type displays the following entry fields and buttons in the Define User Indexes window:



Enter the maximum number of characters you want the indexing oper-1 ator to enter (the system maximum is 239; the default length is one character).

2 Select Yes if you want the index value automatically converted to uppercase before the system stores its values. (You cannot change this selection once you have created a string index.)


Note

Once you have created an index, you cannot delete it or change its type.

If you select No, the index value becomes case-sensitive (each charac-3 ter in the query must match the case of the stored value to qualify for the report).

Numeric Index

Selecting the Numeric index type displays the following entry fields and buttons in the Define User Indexes window:



As an option, you may assign an output mask to a numeric index type. This specifies how the system displays the value (in the Query Match Report, for example). If you do not specify an output mask, the system uses the numeric mask defined in the Image Services server's operating system parameters.

The system accepts only numerals, a minus sign, and a decimal point in numeric fields. Whether document entry operators need to key the decimal point depends on the value and the mask. The mask also determines the index length.

On an Image Services server configured with an Oracle RDBMS, the system does not use the mask size, precision, or scale to validate data entered into the field.

CAUTION

On a Windows Server platform with a SQL database or on all platforms with a DB2 database, the system checks the numeric input value against the precision and scale values defined for a numeric mask on the server. If the input value does not fit into the numeric mask, document committal fails, displaying the error:

<90.0.104> Precision and scale specified in numeric index mask cause overflow

The precision is the total number of digits, on both the left and right side of the decimal point (excluding the decimal point, commas, etc.). The scale is the number of digits shown on the right side of the decimal point. You can enlarge the precision and scale of a numeric mask using the enlarge_ncol command. For details on how to enlarge these values, see "enlarge ncol" on page 496.

Numeric masks can use these characters:

Place holder for a numeral. Specify as many as needed for the # maximum number of digits. Unused positions are blank and the number is right justified.

- Position of a numeral with leading zeros. A digit always displays 0 in each position. Any position for which there is no other value displays a zero.
- Displays the sign (+ or –) of the number. You can place it at the beginning or end of the mask depending on where you want the sign to display.
- Displays the sign of the number only if it is negative. You can place it at the beginning or end of the mask depending on where you want the sign to display.
- \$ Places a floating dollar sign in front of the number.
- Position of a decimal point. A mask ending in .00 indicates fixed point—two digits are always to the right of the decimal and operators do not need to enter the zeros. A mask ending in .## indicates floating point—up to two digits can be to the right of the decimal.
- Position of an embedded comma within the mask to make the value more readable.


The following table shows several examples of edit masks, stored (keyed) values, and their resultant output.

Sample Edit Mask	Stored Value	Sample Display
####	29	29
0000	29	0029
\$####	29	\$29
\$###,###.00	.29	\$.29
\$###,###.00	29	\$29.00

Sample Edit Mask	Stored Value	Sample Display
\$###,###.##	29	\$29.00
+####	29	+29
+####	-29	-29
-####	29	29
-####	-29	-29
####—	-29	29–
-\$###.00	-29	-\$29.00
\$-###.00	-29	\$-29.00
#,###	2929	2,929
##.##	29	29.

Date Index

Selecting the Date index type displays the following entry fields and buttons in the Define User Indexes window:



Enter the mask you want the system to use when displaying this date. Date masks use the following codes:

Day of the week (0–6, where 0=Sunday and 6=Saturday) W

Day of the month (1-31) dd

ddd Day of the year (1–366)

Abbreviated day name (Sun-Sat) day

daynameDay (Sunday-Saturday)

Number of month (1–12) mm

Abbreviated month name (Jan-Dec) mon

month Month (January-December)

Last two digits of year (00-99) уу

Year (0000-9999) уууу


Your mask can include spaces and punctuation characters as separators. Below are some sample date masks and their resulting display.

Date Mask	Display
dayname, month dd, yyyy	Friday, November 10, 1996
dd mon yyyy	10 Nov 1996
mm/dd/yy	11/10/96
ddd	315

For information about how the Image Services system interprets date masks, see "Appendix B - Date and Time Formats" on page 596.

Menu Index

Selecting the Menu index type displays the following entry fields and buttons in the Define User Indexes window:



Click the List button to display a list of existing menus.

If you need to create a new menu, select Build Menu from the Options menu. See "Create a Menu Index" on page 149.

Create a Cluster Index

Clustering is associated with a retrieval key index. You indicate that all documents with the same value in a specified index should be stored on the same storage media regardless of when you scan and commit them. A selected retrieval key used for clustering must already exist.


You can create as many cluster indexes as you need.


Note

You cannot use clustering if your system is configured to use fast batch committal (see "Fast Batch Committal" on page 67).

To create a cluster index:

1 Choose Define/Update Cluster from the Indexes menu to display the Define Cluster Index subwindow.





Click the Add button to display the following dialog box. 2

- 3 Choose a string or numeric index for clustering from the Key Index Name pull-down list.
- From the Family Name pulldown list, choose a media family to write to 4 when using this cluster index.
- 5 Choose one of the two radio buttons to select Average Documents Per Cluster or Reserve One Surface Per Cluster.
 - If you choose Average Documents Per Cluster, enter a number in the text field that represents the number of documents you expect to be in one cluster (maximum is 500).

CAUTION

Average Documents Per Cluster is critical to your system performance. If you set this value too small, your clusters will eventually split among multiple media surfaces, reducing clustering's effectiveness. If you set the value too large, you can waste a significant portion of each surface.

If you choose Reserve One Surface Per Cluster, one entire surface of the media is reserved for the cluster.

- 6 Click OK to accept your settings and return to the Define Cluster Index window; click Cancel to close the Add Cluster Index dialog box without making any changes.
- Choose Save from the File menu and follow the prompts to save your 7 changes.

Assign a Cluster to a Document Class

Once the cluster index exists, you can set up a document class to use clustering. See "Create Document Classes" on page 154. Before you can choose the Cluster Family checkbox in the Define/Update Document Classes window, you must enter the cluster index name at the bottom of the window. Then, when you click the Cluster Family button, the system automatically enters the media family name. Complete the rest of the fields as appropriate and save the document class by choosing Save from the File menu.

Modify a Cluster Index

To modify a cluster index:

- 1 Choose Define/Update Cluster from the Indexes menu.
- 2 Select a string or numeric index and click the Modify button.
- 3 Follow steps 3 through 7 under "Create a Cluster Index" on page 138.

Delete a Cluster Index

To remove the clustering attribute from an index:

- 1 Choose Define/Update Cluster from the Indexes menu.
- 2 Select a string or numeric index and click the Delete button.
- 3 Confirm your choice in the dialog box that appears.

Change Retrieval Key Status

You can change a retrieval key index to an informational index and change an informational index to a retrieval key. These functions are on the Database Maintenance window's Indexes menu.

CAUTION

It is highly recommended that the Define/Update option under the Indexes menu be executed during non-production hours. This tool will issue locks on the relational database (Oracle, SQL Server or DB2) that may cause long delays in system processing. This primarily affects systems with very large databases, but is the recommended procedure for all systems.

The Report function under the Indexes menu provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.


Change a Retrieval Key Index to an Informational Index

When you change a retrieval key index to an informational index, the system searches sequentially when performing a retrieval based on this index and the index can no longer be used as a primary key.

However, the document entry process is faster for subsequent indexing of documents that use this index because the system is not maintaining a sorted list.

Using this option can save magnetic disk space by eliminating the overhead structure required by retrieval indexes.

- 1 Restart the FileNet software.
- 2 Select Drop Retrieval Key from the Indexes menu to display this window.



- 3 Click the down arrow and select the retrieval index from the list.
 - The description and the index type display.
- 4 Select Drop from the File menu and answer the prompt.


Change an Informational Index to a Retrieval Key Index

You can change an informational index to a retrieval key index.

CAUTION

The time it takes to build a retrieval key depends on the number of documents stored in the database. For example, it may take approximately two minutes for every one million documents. You can only build a retrieval key when no one needs to use the database (for example, to commit, retrieve, or print any documents).

- Restart the FileNet software. 1
- 2 Select Build Retrieval Key from the Indexes menu to display this window.



- 3 Click the down arrow and select the index from the list.
- 4 Select Build from the File menu and answer the prompt in the popup window

Modify an Index

Before modifying an index, consider the consequences to all document classes that use the index. Changing a description or adding items to menus does not have a significant impact. However, verify that a change that makes sense in one document class does not adversely affect a different document class. Check the Document Class Report (choose Report from the Classes menu) to see which classes use the index.

CAUTION

It is highly recommended that the Define/Update option under the Indexes menu be executed during non-production hours. This tool will issue locks on the relational database (Oracle, SQL Server or DB2) that may cause long delays in system processing. This primarily affects systems with very large databases, but is the recommended procedure for all systems.

The Report function under the Indexes menu provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.

CAUTION

Any modifications made to the Index fields, Document Classes and Disk Families will not take effect until the FileNet software is recycled.

CAUTION

Reducing the maximum string index length generates a database error when you query on documents that were committed using the larger maximum.

The following table lists the fields you can modify for each index type.

Index Type	Fields You Can Modify
Numeric	description, mask
String	description, maximum length of string
Date	description, mask
Menu	description, name of menu

Note

You cannot delete or change the type of an existing index. You also cannot change the "Convert to Upper Case Letters" setting on a String index.

Although you cannot create new indexes when others are using the system, you can modify them at any time.

To modify an index:

- 1 Select Define/Update Index from the Indexes menu in the Database Maintenance window.
- 2 Click the List button to display the list of indexes.
- 3 Select the index to be updated and make changes as needed.
- 4 Save the changes by choosing Save from the File menu.

Rename an Index

If you must rename an index, do it very early in the development process and only if you understand the implications for the entire system. For example, if you are not using an index, you can rename and use it instead of creating a new index (if the index is the correct type). If you change the name of an index that is in use, you may need to make other changes, too. Discuss this change with your service representative.

CAUTION

It is highly recommended that the Define/Update option under the Indexes menu be executed during non-production hours. This tool will issue locks on the relational database (Oracle, SQL Server or DB2) that may cause long delays in system processing. This primarily affects systems with very large databases, but is the recommended procedure for all systems.

The Report function under the Indexes menu provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.

CAUTION


Any modifications made to the Index fields, Document Classes and Disk Families will not take effect until the FileNet software is recycled.

CAUTION

In a Filenet P8 Content Federation Service for IS environment, do not change the name of any mapped index unless absolutely necessary. If you do need to change the index name, you must remap the index on the Content Engine system. Otherwise, index properties exported from IS will be mismatched with Content Engine mapping.

To rename an index:

1 Choose Rename from the Database Maintenance window's Indexes menu to display this window:



2 Select the index you want to rename from the pulldown list.

The Description and Type fields display and the cursor moves to the New Index Name text field.


- 3 Type the new name.
 - See "Indexes" on page 91 for naming conventions.
- 4 Select Save from the File menu. A dialog box asks for confirmation.
- 5 Click OK to confirm or Cancel to exit without saving.
 - If anyone used the database between the time you restarted the FileNet software and the time you saved the name change, you see a message asking you to try later.
- Rebuild indexing forms that use this index. (For details, see your 6 FileNet client software documentation.)

Create a Menu Index

To create a menu index:

Choose Build Menu from the Options menu of the Define User Indexes 1 dialog box (see also "Menu Index" on page 137).

The following window displays.



2 Click the List button to display a list of menus and the blank text field beneath the list.

3 Enter from 1 to 14 alphanumeric characters as a menu name.

Do not imitate batch names (a, b, d, f, q, p, pv, t, w, x, or X followed by a number) or use Visual Workflo reserved words. You cannot start a name with F_ (capital F, underscore) or with a numeral, but you can use an underscore as part of the name.


Click OK. 4

> The selection box disappears and the menu name appears in the Name field.

5 Enter a description (optional) for the menu.

> A description can be up to 177 characters. Only the first 55 characters display without scrolling.

Click the Add button and the Choice Attributes dialog box appears: 6



In the Return Code field, enter a character to be used as a return code. 7

A return code is a single character that tells the system (including a WorkFlo script) which menu option an operator selected. Specify any character, as long as each return code in the menu is different.

Note

A return code can be any upper- or lower-case letter in the alphabet, a number (0-9) or a keyboard symbol. 94 different return codes are available.

When the indexing operator selects an item from the menu, the system passes the return code to Visual Workflo. The item does not appear on the display during indexing. If retrieval operators use a menu index as the basis for sorting the Query Match Report, the return code determines the sort order. To avoid confusion, you might want to assign letters of the alphabet to sort the menu items alphabetically.

Note

The IDM Desktop Find program cannot locate a document based on a single-digit menu item selection unless its return code matches the text displayed on the menu. If you are defining a single-digit menu item, you must assign it an identical return code. For example, if you create a menu item in the Text to Display field as an upper-case letter A, you must also enter an upper-case letter A in the Return Code field.

8 In the Text to Display field, enter text for a menu item.

After entering the return code, press Tab and enter the text you want to appear on the menu for the indexing operator.

9 Click the OK button.

> The new menu item is added. You can continue to add entries, clicking OK after each one.

- When you finish adding entries, click Cancel to go back to the Build 10 Menus window.
- 11 Choose Save from the File menu and a dialog box pops up to tell you that the menu was successfully created.
- 12 Exit the dialog box by choosing Exit from the File menu.

Change a Menu

You can add, modify, or delete items in existing menus, or you can copy, rename, or delete menus.

CAUTION

It is highly recommended that the Define/Update option provided for the Document Classes be executed during non-production hours. This tool will issue locks on the relational database (Oracle, SQL Server or DB2) that may cause long delays in system processing. This primarily affects systems with very large databases but is the recommended procedure for all systems.

The Report function under the Classes menu provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.

CAUTION

Any modifications made to the Index fields, Document Classes and Disk Families will not take effect until the Image Services software is recycled.

Add, Modify, Delete Items in a Menu

- 1 Choose Define/Update Index from the Indexes menu.
- Choose Build Menu from the Options menu. 2

- 3 Select the menu you want to change by clicking the List button and selecting an item from the list box. Click OK.
 - To add an item to the menu, click the Add button.
 - To modify an item in the menu, select the item and click the Modify button.

In each case, enter the appropriate text in the resulting dialog box and click OK.

- To delete an item, select it and click Delete. No confirmation prompt appears.
- Save your changes by choosing Save from the File menu. 4
- 5 Select Exit from the File menu.

Copy, Rename, Delete Menu

- Choose Define/Update Index from the Indexes menu. 1
- 2 Choose Build Menu from the Options menu.
- 3 Select the menu you want to copy, rename, or delete by clicking the List button and selecting an item from the list box. Click OK.
- From the File menu, choose the appropriate action: 4
 - Save As to make a copy of the menu using a different name.
 - Rename to give the menu a different name.
 - Delete to delete the menu entirely.

Save As and Rename present a prompt for a new name. Enter the new

Delete asks you to confirm that you want to delete the item. To delete the menu, click Yes.

5 Select Exit from the File menu.

menu name and click OK.

Create Document Classes

Before you create a document class, the indexes and media families needed by the class must already exist.

Note

In a FileNet P8 Content Federation Services environment, a document class may or may not have index values associated with it.


Define only as many document classes as you need. You cannot delete a document class once you define it.

Regardless of whether you are adding or modifying a document class, the change takes about five minutes to be effective at the workstations.


CAUTION

Do not modify a document class if any uncommitted batches exist that use that document class. These batches cannot be committed if the document class changes.

1 Choose Define/Update Class from the Classes menu to display the following window.



- 2 Define document class.
 - а Click the List button to display a dialog box listing document class names you can modify.



- To modify an existing document class, choose the name of the document class you want to modify and click OK.
- To create a new document class, type the new name in the Edit field at the bottom of this list and click OK.

Note

In a FileNet P8 Content Federation Services environment, select a name such as CFS_apps or P8_acct, that will be easy to identify when you map it to the Content Engine document class.

Enter a description.

You can enter up to 30 characters (of any type or case) that describe the purpose of this document class. The description appears on the document class report and is required. If you chose an existing document class to modify, the system automatically enters the description. You can change the description or leave it as is.

Select the media family name.

If you chose an existing document class to modify, the system automatically enters the media family name. Use the pull-down list to select a primary media family name. If documents in this class will not be stored on storage media, you must still enter a media family name.

Important!

When you make any configuration change to Image Services through the Application Executive (Xapex) such as reassigning a media family or adding an index, you must restart any associated client server that runs IS Toolkit-based applications. This includes CFS-IS servers in a FileNet P8 Content Federation Services environment. For example, If you assign a mapped IS Doc Class to a different media family in a FileNet P8 Content Federation Services environment in New York, you must restart the associated Content Engine server in Chicago. Otherwise, the Chicago server will not recognize the change and may commit documents to the wrong media.


If you store documents in clusters, check the Cluster box.

You must also enter the currently defined cluster index in your index list (see "Assign a Cluster to a Document Class" on page 141). Talk with your service representative before choosing this option. This setting is ignored if your system is configured to use fast batch committal (see "Fast Batch Committal" on page 67).

- 3 Define DMA properties. The DMA Properties enable you to specify a DMA display name and up to 10 GUIDs (Globally Unique IDentifiers) for each defined document class.
 - Specify the Display Name. When you first enter a name in the Document Class Name field, the system automatically duplicates that name in the DMA Properties Display Name field. You can change this name by selecting all or part of the text and typing the desired name.

This field accepts up to 30 alphanumeric characters, including special characters. The following display name restrictions apply:

- space, tab, colon (:) and equal sign (=) are not allowed.
- cannot start with a single or double quotation mark.
- cannot end with a backslash (\) character.
- b Click the Edit button to open the GUIDS dialog box, which enables you to add, delete, or modify GUIDs for the currently defined document class. You must assign at least one GUID to the document class and may assign up to 10 GUIDs for each document class.



- Select the GUID you want associated with the document class. For details on adding, editing, and deleting GUIDs, see "GUIDs Assigned to Document Class" on page 167.
- d Click OK. The selected GUID displays in the GUIDS field on the Define/Update Document Classes window. Each additional GUID defined on the GUID list may be an alias for the document class.
- 4 Click the Yes radio button to enable cataloging (the default).

To disable cataloging for a document class, click the No button. To reenable cataloging, click the Yes button. The system keeps track of when you turn cataloging on and off via a table in the index database.

CAUTION

If you turn off cataloging and commit documents to the document class, the index information will not be in the index database, but the address information will still be in the permanent database, causing a discrepancy between the databases. Import the documents from the media to get the index database information (see "Importing Documents" on page 374).

- 5 Specify migration options.
 - Select a migration option (click Yes or No).

Select Yes to migrate your documents to storage media. This setting is ignored if your system does not use storage media or if you are doing a fast batch committal (see "Fast Batch Committal" on page 67).

Note


If you have a cache-only, multi-server system and you have configured phantom storage libraries, set "Migration to O.D." to No.

Set migration delay.

Click Yes to delay migration, then set the number of days and hours until migration starts. At the end of this time, documents migrate to storage media and the space they occupy in page cache can be reused. The maximum delay you can set is 24,855 days and 3 hours.

Set security options. 6

Select Change Access from the Security menu to display this dialog.



- To restrict any of these functions (read, write, or append/execute) to one user or group, replace the name (ANYONE) with the name of the user or group that should have each kind of access.
- To change access back to anyone, you must type (ANYONE), including the parentheses and using uppercase letters.
- 7 Set document entry parameters.
 - Set pages per document.
 - If the number of pages usually varies, leave the Variable button selected.
 - If documents in this class usually have the same number of pages, click the Fixed button. A text field appears where you enter the number of pages (up to three digits). This number serves as a default for the operator who defines the batch, but the operator can change the default as required for any one batch.
 - b Enter maximum pages per batch.

Click in the text field and type the maximum number of pages you expect to scan in one batch.

Set verification options. 8

Click the Image, Index, and Batch Total checkboxes as required. You can select any combination, as long as you use indexes with the appropriate definition.

Verification Options	Description
Image	Check this box to require the operator to verify the scanned image on screen as an extra step between the scanning and indexing steps. Even if you do not select this option, operators can verify images for any batch at any time.
Index	Check this box to set up a verification pass as a default step to be performed for documents in a class. PC workstation operators can verify indexes at any time.
Batch Total	Check this box and the system adds up the values of a numeric index. After the batch is completely indexed, the total for that numeric index should match a predefined batch total entered during the define batch session.
	Batch total verification also verifies the number of pages and documents expected in the batch. You can enable batch totals as a default, or as an optional step for any particular batch, only if the document class includes a numeric index enabled for batch totals. Document entry operators can perform batch total verification only if you set up the document class to use it.

9 Set disposition parameters.

In this section of the dialog, you indicate how you want to eventually dispose of the information stored in the index database for documents in this class. You also indicate how long to wait before flagging the information and when to start counting the time.

Disposition Parameter	Description
Archive	Choose Archive to move the information to an archive database on storage media. You can then query the archive database and still access documents, even though the index information is no longer in the index database on magnetic disk.
Delete	Choose Delete to remove the information from the index database permanently and never access the documents on storage media again. If you use erasable storage media and plan to reuse the media, you would most likely choose this option.
Months From	In this text field, indicate how many months you want to keep the index information available in the index database. The actual time index information remains in the index database depends on when you start counting the time (committal date or closing date) and when you run the processes that actually archive or delete the information.
Date Filed	Choose Date Filed to start counting time on the day the system writes the document information to the index database.
Date Closed	Choose Date Closed to start counting time when an operator closes the document.

Route data about scanned documents to a Visual Workflo queue 10 (optional).

Parameter	Description
System	Click the arrow to the right of the System field and select the name of a workspace from the pulldown list.
Queue	Click the arrow to the right of the Queue field and select the name of a queue from the pulldown list.


To delete information about scanned documents from a Visual Workflo. queue:

- To remove all the gueues in a system, double-click on a system name in the list (not in the System field, as this has no effect).
- To remove only one queue in a system, double-click on a queue in the list (not in the Queue field, as this has no effect).
- 11 Specify the indexes.


Note

In a FileNet P8 Content Federation Services environment, a document class may or may not have index values associated with it.

To assign indexes to the document class, choose Edit from the Index menu. The following dialog displays.



Click the Add button to specify the indexes that belong to this document class. The following dialog displays.



Click the arrow to the right of the Name field to display a list of indexes. Click a name to select it.

Select the document entry options for the index.

Document Entry Option	Description
Verify	Click the Yes button to default to index verification for this index. When indexing accuracy is important, you can require a second entry operator to verify index values. Verification consists of indexing the batch a second time and comparing the results.
Batch Total	For numeric indexes only, click the Yes button to collect totals. The system also keeps track of the number of pages and documents scanned. You must select this option as a default for the class before document entry operators can select it when defining the batch. If the expected totals entered during batch definition do not match the totals calculated by the system during document entry, the entry operator can resolve the inconsistency before committing the batch.
Required	Click the Yes button to require an entry for this index. If the indexing operator can leave the field blank, leave the default of No selected.
Auto Indexing	For string indexes only. The client program used for scanning determines the setting for this value:
	WorkFlo/Scan is your scan client. Click Yes if you have some means of acquiring index data automatically (bar codes, patch codes, etc.). WorkFlo/Scan can perform automatic indexing only if you specify Yes in this Auto Indexing field on the Image Services server.
	Capture Professional is your scan client. Click No. The Capture settings collection contains the automatic indexing configuration. Capture ignores the Auto Indexing setting on the Image Services server.

- When you are satisfied with the specifications for an index, click the OK button at the bottom of the dialog box. The information is transferred to the User Defined Index dialog box (though it may be hidden behind the currently displayed dialog box).
- To define another index, select another name. To finish selecting indexes for the document class, click the Cancel button.
- Click Close on the User Defined Index dialog box. f

12 Save the document class.


Select Save from the Define/Update Document Classes window's File menu. A message box appears, telling you when the document class is successfully created. You must acknowledge the message by clicking the OK button.

GUIDs Assigned to Document Class

GUIDs (Globally Unique IDentifiers) are DMA-compliant, 16-byte integers used to uniquely identify each element transported over a network. The system ensures unique GUID assignments by automatically generating this integer using an algorithm based on the system's network card MAC address and a format that complies with the specifications provided for the system's platform.

Each GUID must have a unique name that conforms to the format specified for your Image Services system platform. For platform specifications, see the GUID naming conventions described in the documentation that came with your operating system.

To add, delete, or rename GUIDs for the current document class, click the Edit button in the DMA Properties box of the Define/Update Document Classes window. The GUIDs dialog box opens, displaying a list of currently defined GUIDs.



As you add or rename GUIDs, the system validates the assigned name, making sure it meets the following requirements:

- Unique name assignment
- No duplication of name in the Image Services database
- At least one and no more than 10 GUIDs assigned to each document class


You may modify the currently defined list as follows:

To delete an existing GUID, select it from the list and click the Delete button. A message box opens, asking you to verify deletion. Click Yes to delete the selected GUID or No to cancel deletion.

- To add a new GUID to this user index, click the Add button. The Add New GUIDS dialog box opens. For details, see "Add New GUIDS" on page 169
- To rename an existing GUID, select it from the list and click the Edit button. The Edit GUID dialog box opens, displaying the exiting name of the selected GUID. For details, see "Rename GUID" on page 170.

Add New GUIDS

The Add New GUIDS dialog box opens after you click the Add button on the GUIDS dialog box when defining the DMA properties of a document class.



Automatically Generate GUID

Click this radio button to have the system automatically generate a GUID. The system displays the GUID in grayed-out, uneditable text. To save the automatically generated GUID, click the Ok button. The system closes the Add New GUIDS dialog box and displays the newly assigned GUID in the GUIDS list.

Manually Enter GUID


Click this radio button if you want to manually assign a GUID. Type a 16-byte character set into the field, using the format shown in the sample displayed below the entry field.

To save the manually defined GUID, click the Ok button.

- If you entered a GUID that already exists in the Image Services database or does not use the correct format, a message box opens, alerting you to the problem. When this happens, you must enter an acceptable GUID and click Ok again.
- If you entered a unique GUID using the correct format, the system closes the Add New GUIDS dialog box and displays the newly defined GUID on the GUIDS list.

Rename GUID

The Edit GUID dialog box opens after you select an existing GUID from the GUIDS list and click the Edit button on the GUIDS dialog box when defining the DMA properties of a document class.



Edit existing GUID

This field displays the currently selected GUID and allows you to rename it.

To change the current name, select all or any part of the displayed text, type in the desired text, and click Ok.

- If you entered a GUID that already exists in the Image Services database or does not use the correct format, a message box opens, alerting you to the problem. When this happens, you must enter an acceptable GUID and click Ok again.
- If you entered a unique GUID using the correct format, the system closes the Edit GUID dialog box and displays the renamed GUID on the GUIDS list.

Modify a Document Class

You can modify a document class. Most changes are straightforward. You simply enter different information in the same manner as creating a new document class.

CAUTION

It is highly recommended that the Define/Update option provided for the Document Classes be executed during non-production hours. This tool will issue locks on the relational database (Oracle, SQL Server, or DB2) that may cause long delays in system processing. This primarily affects systems with very large databases but is the recommended procedure for all systems.

The Report function under the Classes menu provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.

CAUTION

Any modifications made to the Index fields, Document Classes and Disk Families will not take effect until the Image Services software is recycled.

Be aware of how changing a document class affects your system. You might need to complete several steps, including updating custom indexing forms on PC workstations if you add or delete an index. If you do not complete these steps, indexing cannot function correctly.

If you add an index to a document class, you can only use the index to retrieve documents committed after the addition. If you delete an index from a document class, you cannot retrieve any documents based on that index. You should plan carefully and test your document classes thoroughly.

CAUTION

Be sure that **no uncommitted batches** use the document class you are modifying. Otherwise, operators must repeat the work for the batch to commit successfully.

Change Indexes

To change an index:

- 1 Choose Define/Update Document Class from the Database Maintenance Classes menu.
- 2 Select the document class from the List menu.
- 3 Choose Edit from the Index menu.
 - To modify an index, select it from the User Defined Index list and click Modify. In the resulting dialog box, change items by clicking the Yes or No button. If changing autoindexing information, you can directly alter the text in the grid.

- To delete an index, select the index and click the Delete button. You must rebuild the custom indexing forms used at PC workstations if you delete an index.
- To add an index, click the Add button and define the index using the procedures described under "Define an Index" on page 119. You must rebuild the indexing forms if you add an index.
- 4 Click the OK button.
- 5 When you finish modifying all indexes, click Close to return to the User Defined Index dialog box.

Change Security

To change security for a document class:

- 1 Select Change Access from the Security menu.
- 2 Enter the new names of groups you want to authorize to read, write, or perform other operations on documents in this class.

When you change security for the document class, only new documents scanned into the class acquire these settings. To change all documents previously scanned into the class, select Update Doc. Security from the Miscellaneous menu (see "Update Document Security" on page 185).

Save Changes

Select Save from the File menu to save any changes.

Update Retention Parameters

With Update Retention Parameters (from the Miscellaneous menu), you can change the retention parameters of documents and folders. You must be logged on to the root/index server and be a SysAdmin or user with full administrative rights (Admin Group, Session Group, Primary Group and Member of Group are all set to SysAdminG) to use this option. Update parameters only on the root/index server and only when no other activity is taking place on the system.

CAUTION

It is highly recommended that the Define/Update option provided for the Document Classes be executed during non-production hours. This tool will issue locks on the relational database (Oracle, SQL Server, or DB2) that may cause long delays in system processing. This primarily affects systems with very large databases but is the recommended procedure for all systems.

The Report function under the Classes menu provides a method to list these definitions and should be used in place of the Define/Update option whenever possible.

CAUTION

Any modifications made to the Index fields, Document Classes and Disk Families will not take effect until the Image Services software is recycled.

When you select Update Retention Parameters from the Database Maintenance window's Miscellaneous menu, the following window appears.

Note

On a Remote Admin Console (RAC) system whose logon is remote, this option will be disabled (grayed out). For more information about RAC, see Remote Admin Console User's Guide. To download this

guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.



In the top section, you select documents and folders to be affected by the change. The way you select documents (by document class or by document ID) determines how you can change retention parameters. In the bottom section, you select the types of changes to make. Regardless of the retention parameters you select, the process always corrects invalid entry dates for selected documents.

CAUTION

Updating retention parameters can change the dates associated with every document on the system. Use care in selecting what you want to change, especially when specifying an entry date.

Prior to archiving or deleting documents from the index and permanent database, review retention parameters and make any desired changes (see the table under "Choose Retention Parameters" on page 183).

For example, you may decide to archive documents that are currently set to be deleted. You could select one or more document classes and set the Dispose Type retention parameter to Archive. If you had previously set the document class to be archived, then you could just execute Update Retention Parameters without changing any retention parameters. This would update all documents according to the current retention parameters in the class. Updating retention parameters corrects any invalid dates (such as documents with a date of entry earlier than the entry date of the lowest numbered active document). Updating retention parameters also corrects any invalid document dates. It does not validate folder dates. Be aware that even if no parameters are selected, but documents are, the entry date may change.

You can either supply a date to be used as the earliest acceptable date or use the default date (the entry date of the lowest-numbered active document on the system).

If you plan to use the default date, you should perform a guery and check the date of the lowest numbered document before updating retention parameters. Any changes made, individually, to documents and folders will be lost when those documents and folders meet your selection criteria. Also be aware that the entry date will change even if the Entry Date check box remains unchecked.


Select by Document Class

Selecting documents by document class assures that all documents in a class have the same retention parameters. If you enter retention parameters, the process changes all documents in the class. If you do not enter retention parameters, the process only checks for invalid dates and updates documents to match retention parameters defined for the document class.

To select by document class:

- Check the Classes radio button. 1
- 2 Check the All box and the Selected List displays all document classes.

All documents on the system are updated according to the current settings of the document class or the retention parameters you choose at the bottom of the window. Be sure any new retention parameters are appropriate for every document on the system.



Check Select From List for a scrolling list box of all document classes. 3

- Selecting one or more of these classes (click or drag through them), makes the Add button available.
- Click Add to add the highlighted classes to the Selected List.
- Delete items by choosing one or more names in the selected list, then clicking the Delete button.

Select by Document Number


Selecting documents by document number applies retention parameters to all documents in the range you specify. This method of document selection does not look at the document class.

If you use only one document class, or all of your classes use the same retention parameters, this method of selection lets you specify as many documents as you have time to update in one run. When that group is finished, you can start with the next higher document ID and update another group of documents.

If you need to update large numbers of documents, this is an efficient way to proceed. You might do several hundred thousand one weekend and several hundred thousand the following weekend.

To select by document number:

- 1 Click the Documents radio button.
 - If you choose All documents, no further input is required.
 - Choose retention parameters as needed, but make sure they are appropriate for all documents on the system.
- 2 If you choose Build Range/List, you must enter starting and ending document numbers (the lowest acceptable document number is 100,000).



Select by Folders


Selecting by folders changes folders with older dates or dates later than today's date to the entry date you specify. This is not an alternate method of selecting documents filed in folders.

Selecting by folders does not check for invalid folder dates unless you specify an entry date (no default date exists for folders). When you use an entry date, you are specifying that it is the oldest valid date.

To update the retention parameters of folders:

- Choose the Folders radio button. 1
- 2 Choose All to have all folders on the system changed to the retention parameters you choose.
- 3 Choose Select From List to display actual folder names that meet your qualifications.

The display limit is 100 folder names, beginning with the lowest entry. Names are case sensitive and always start with a slash (I). You do not have to enter complete names. For example, entering /a selects all folder names that begin with a lowercase a.




Enter the earliest alphabetic entry in the Starting Folder Name field. Enter the last folder name to be updated in the Ending Folder Name field and click OK. A message box displays, telling you how many qualifying folders were found, along with the last (100th) folder name.

- b Click OK and the main dialog box shows an empty Selected List scrollbox, an Available List scrollbox with the folder names that qualified, and grayed out Add and Delete buttons.
- Select folders from the Available List and click Add to add them to the Selected List. The Selected List then becomes the list of folders you want to update.

If the range you specified includes more than 100 folders, a message box informs you of that and tells you the name of the 100th folder. You can specify that name (or the next name) as the starting name when you run the process again for the next 100 folders.

Choose Build Range/List and click the Add button in the Selected List 4 scrollbox to build a list of folders.



In the dialog box that appears, enter part of a folder name with a percent sign (%) at the end to signify all remaining characters. Clicking OK adds this specification to the Selected List. You can specify more than one wild card name. You can select as many folders as you wish by this method (you do not see the list of selected folders).

Choose Retention Parameters

The four retention parameters are Entry Date, Base Type, Offset (month), and Dispose Type. If you select by document class, you do not have to choose retention parameters for documents. If you have already modified the document class, all documents in the class will be changed to match the current retention parameters. When selecting by document or folder and you want to change something other than invalid document dates, you must choose the appropriate retention parameters. You may want to select Update Retention Parameters just to check for invalid dates (see "Update Retention Parameters" on page 174).

If you don't enter an Entry Date, the system uses the entry date of the active document with the lowest ID. Any entry date earlier than that or later than the current day's date is considered invalid and will be changed to the entry date of the active document with the lowest ID (documents only). If the entry date changes, the process also recalculates archive and delete dates when they are based on entry date.

For example, suppose that on 3/31/2001 you have these documents:

Doc ID	Entry Date	Valid/Invalid
1001	2/25/2001	Valid
1002	2/26/2001	Valid
1003	2/27/2001	Valid
1004	2/28/2001	Valid
1005	2/17/2001	Invalid (Earlier than the lowest numbered document)
1006	3/15/2001	Valid
1007	4/15/2001	Invalid (Later than today's date - 3/31/2003)

- If you don't enter an Entry Date, only the two documents whose dates are invalid will be changed to the entry date of the document with the lowest Doc ID, in this case, 2/25/2001.
- If you do enter an Entry Date, the documents with entry dates earlier than this will be changed to the date you just entered.

CAUTION

Do not enter today's date! If you enter the current day's date, the system handles ALL earlier dates as invalid and updates them.

The following table describes the retention parameters.

Retention Parameter	Description
Entry Date	The earliest possible entry date. This corresponds to the first documents you scanned. The purpose of this parameter is to identify invalid dates. To enter a date, select the place holder text (yyyy/mm/dd) and replace it with valid numerals using that format.
Base Type	Entry means date of entry (date cataloged). Closed means date an operator closed the document or folder.
Offset (month)	Number of months to wait after the base event occurs. If the base is date of entry, this will be the total time you want the document to remain in the database. If the base is date closed, the document remains active until an operator closes the document. In either case, when the time expires, the document becomes eligible for removal, either by archiving or by deleting.
Dispose Type	Choosing Delete permanently removes the document or folder information from the databases when you run Delete Doc/Folder from Database Maintenance. (On cache-only systems, this also deletes the document from cache.) Archive moves the document information from the databases on magnetic disk to archive databases on storage media.

Start the Update

After choosing what to modify and how to modify it, click the Execute button to begin processing. A progress indicator shows the number of records found so far and the last record updated. A checkpoint file is also created to show the progress:

/fnsw/local/tmp/upretent.CHK (UNIX)

\FNSW LOC\tmp\upretent.CHK (Windows Server)

When the process completes, a dialog box shows you how many documents or folders qualified and how many were updated.

Update Document Security

In a system with a storage library or an ODU, select Update Document Security to update the security information for previously committed documents. If you change the access specifications for a document class after using the class to commit documents, older documents in the same class will have different group access than newer documents. Update Document Security makes the group access the same for all documents in the class.

Note

On a Remote Admin Console (RAC) system whose logon is remote, this option will be disabled (grayed out).

To use this option, you must be logged on to the document locator (primary storage library) server as SysAdmin or as a user with full administrative rights (Admin Group, Session Group, Primary Group and Member of Group are all set to SysAdminG).


Updating document security can take a long time if you specify a large number of documents to update. The process looks at every document in the document number range you specify for each document class you specify. However, you can stop the process and restart it later without having to start over from the beginning.

CAUTION

Do not run any other process that accesses the index database while updating document security. In general, this means that no other users should be accessing the Image Services server.

To update document security:

- Select Update Doc. Security from the Miscellaneous menu. 1
- 2 Click Yes at the warning prompt to display the Update Document Security window.



- 3 Enter the lowest document number in the range of documents you want to update. The default is the first document in the database.
- 4 Enter the highest document number in the range of documents you want to update. The default is the last document in the database.
- 5 Select document classes.
 - To specify all document classes on the system, check the Select All Classes checkbox.
 - To select document class names, click a name in the Available Document Classes list and click the Add button to add the name to the Selected Document Classes list.
 - To remove a class from the selected list, click the class name to highlight it, then click the Delete button.

The process updates only documents belonging to the document classes you specify.

6 Click the Execute button to start the update.

> The process updates all documents within the specified range that belong to the document classes you indicated. It changes the security information for committed documents to match the current definitions in the corresponding document classes.

As the process continues, it displays the names of the document classes and shows you how many updates it has made.

7 Click the Interrupt button if you need to stop the update procedure.

The process checks for a cancel request after every 100 records. It also creates a checkpoint file (/fnsw/local/tmp/Xupsec.CHK on UNIX platforms; \FNSW_LOC\tmp\Xupsec.CHK on Windows Server) showing the number of the last document updated. If you stop the update, use the last number updated as the beginning document number when you restart the update.

Delete Expired Documents and Folders


You must be logged on to the Index Server as SysAdmin or as a user with full administrative rights (Admin Group, Session Group, Primary Group and Member of Group are all set to SysAdminG) to delete documents or folders.

CAUTION

Before you use this process, be absolutely sure you will never want to access the documents again.

To delete documents or folders:

1 Select Delete Doc/Folder from the Miscellaneous menu to display the Folder and Document Maintenance window.



2 Click Folders, Documents, or both, then click OK.

During the deletion, you can see progress messages in the large field at the top of the window. The three small fields show the number of documents removed from folders and how many folders and documents were deleted.

3 To terminate the deletion, choose Exit from the File menu.

The system deletes all documents and folders eligible for deletion. The next time you delete expired documents and folders, the system deletes any additional documents or folders that are now obsolete, as well as those it missed when you terminated the process.

Note

You will not be allowed to delete documents or folders on or before the deletion date assigned to them. You can only delete documents or folders after their assigned deletion date.

Get Database Reports

The Indexes, Classes, and Families menus each include a Report option. Select the report for the kind of information you are interested in: user indexes, document classes, or media families.

Save Reports to a File

To save all or part of the report information in a file, you must first select one or more entries. Either double-click the items you want to include or use the Edit menu's Select All option to select all entries at once. After selecting the appropriate entries, choose Save As from the File menu and follow your operating system's procedure for saving files.

Print a Report

You can print or fax all or part of a report, based on how many items you select. Choose Print from the File menu and complete the fields in the dialog box. (For details, see "Print File or Report Dialog Box" on page 543.)

Find an Index, Document Class, or Media Family


If a large number of indexes, document classes, or media families exist in a report, you can use the report window's Retrieve menu to go to a particular name or number.

Select **Go to Name** to display a text popup where you can enter the name of the index, document class, or media family. That item then appears at the top of the report window. If the names of two or more items start with the same characters, you must type enough characters to distinguish between the names (the names are case sensitive).

Select Go to Number to find the index, document class, or media family by the number in the upper left corner of each row in the report. Indexes, document classes, and media families are numbered in the order they were created.


User Indexes Report

The User Indexes Report shows parameters for each index. The information varies based on the type of index. The display shows as many indexes as fit in the dialog box. Click in the gray area of the scroll bar to scroll to the next page of indexes.



Document Class Report


The Document Class Report lists all document classes on the system, one at a time. The amount of information shown depends on the size of the window.



Use the scroll bar at the right of the window to view all the indexes (scroll up or down).

Media Family Report

The Media Family Report displays the names of all media families on the system and shows how the disks are configured.



Use the scroll bar at the right of the window to view all the media families (scroll up or down).

Security Administration

Overview

Security Administration allows you to control user logons, passwords, devices, and the group memberships that determine access to system data and functions (menu options).

Since the FileNet system, including WorkFlo, ignores any security attached to a file at the operating system level, you must use Security Administration to secure the FileNet system. However, you add or update document security through the Database Maintenance program (see "Update Document Security" on page 185).

You turn printer and fax server security on and off through the Configuration Editor (see the System Configuration Tools online help).

Basic Concepts

Security Administration is governed by a set of basic concepts that apply to a number of different parts of the system. When you understand these concepts, the interactions between the different parts of the Security Administration system are easier to understand.

Security Object

Users, groups, and devices are security objects. You can apply most security characteristics to users, groups, and devices.

User

A user is a security object that can log on and perform tasks.

Group

A group is a security object to which you can assign one or more users, devices, or other groups. Some group assignments confer membership; other group assignments, such as those of administrative or session groups, do not confer membership. A group cannot log on or perform tasks. You define data objects with groups or users, including (ANYONE) or (NONE), with read, write, and append/execute permissions.

Device

The two major types of security devices are:

- **Terminals**
- Printers and fax servers

Terminal security controls logons and data access from the terminal. Printer and fax server security controls access to print and fax devices. Setting up printer and fax security is a two-step process, involving both the Configuration Editor and Security Administration.

Permission

Permission is the privilege granted to a security object to perform certain tasks. Each security object has permissions based on its own definition, as well as the definition of the system and the groups of which the object is a member. The extended memberships and override capabilities of the requesting users, groups, and devices control permissions.

Administrator

An administrator is a user assigned special privileges to perform security tasks affecting users, groups, and devices. The four administrative attributes are: supervisor, principal, group, and password. By assigning various combinations of administrative attributes to different users, you can provide checks and balances in your security system.

Membership

You can assign membership for users and devices in one or more groups. A user or device inherits the permissions of all groups of which the user or device is a member. You can assign a group membership in one or more groups. The group inherits the permissions of all groups of which it is a member.

Extended Membership

A user is a member not only of the groups to which the user is assigned, but also to any groups to which those groups are, in turn, assigned. Use extended membership with caution. Be certain you know who is a member, directly or through extended membership, when you expand privileges by assigning one group to another.

You can use group membership to control who can log on from which terminals, access which data, and use which print and fax devices. A security object can belong to any number of groups, either directly or through extended membership. For example, if a user is a member of group A, and group A is a member of groups B and C, and group C is member of groups D and E, then the user is a member of groups A, B, C, D, and E. Through extended membership, the user inherits the permissions of all these groups and can perform tasks that require permissions beyond those explicitly assigned to group A.

Administrative Domain

Any user with administrative attributes can manage his own administrative group. This is the user's administrative domain.

Session

A session is a single logon occurrence. You can assign users to a session group, allowing addition or modification of logon privileges for an entire group of users at once. Assigning a session group is optional.

Data Object

Documents, folders, and annotations are data objects. You can assign security to these objects through the groups, granting them read, write, and append/execute privileges using the Database Maintenance program (see "Document Security" on page 102).

Override

The system override attribute causes the security object's attributes to be used instead of the system attributes in certain instances. Attributes for which overrides apply are: device security, time use restrictions, security logging, and maximum concurrent sessions.

You can set system overrides for groups, users, and devices. Permission to perform a given task depends on how you set these overrides for the user, the terminal, any other device required, and the extended group membership of the user, terminal, and device.

You can reserve the ability to permit groups, users, or devices to override the system defaults for SysAdmin or grant it to other administrators.

Template

Templates are security objects with default attributes to assist you in creating users, groups, and devices. When you create a user, group, or device, it inherits its attribute values from its template. You can customize the three templates provided by Security Administration.

Expiration

Expiration makes a user, group, or terminal unusable after exceeding the time set by an administrator. You can also set user accounts to an expired status if the user does not renew a password when required or if he has exceeded the permitted failed logon attempts. The system does not delete an expired security object, which the appropriate administrator can reactivate.

Logging

The system writes security logs daily and keeps them for 28 days. On the 29th day, the system overwrites the oldest security log. To view security logs, select Events Log from the System pulldown menu. Note that security logs are different from system event logs.

Security Database

The security database sec db0 is a multi-keyed file (MKF) database that contains security information for:

- the FileNet system
- each object (user, group, device)
- each direct membership occurrence
- each function name and class
- each database logon

The system stores security information with its related data. For example, the system stores document security along with the document indexing information in the index and permanent databases; it stores annotation security in the permanent database.

You must back up sec_db0 along with other MKF databases. The system includes security event logs when you back up other data in the fnsw directory. Refer to the manual provided for your backup program: Enterprise Backup/Restore User's Guide, System Administrator's Companion for UNIX, or System Administrator's Companion for Windows Server. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

System Security

As SysAdmin, you can set the default security for the FileNet system. You can use the system default settings provided with the system or you can modify the default system security. See "Set Up FileNet System Security" on page 228.

Group Security

A group is a security object to which you can assign one or more users, devices, or other groups. A group cannot log on or perform tasks. You create a group and assign security objects to the group for one of two reasons:

- To permit an administrator to set values for several security objects at once
- To grant the object certain privileges held by the group, such as access to devices, data objects, and functions

The system reserves the following groups for special uses:

- (NONE)
- (ANYONE)
- SysAdminG
- AuditG
- FieldServiceG
- OperatorG

For details, see "Reserved Groups" on page 202.

The following group assignments have special meaning:

- **Primary Group Assignment**
- Administrative Group Assignment
- Session Group Assignment

For details, see "Group Assignments" on page 203.

Administrative and session group assignments do not provide membership privileges in the group.

Reserved Groups

The following table describes the six reserved security groups: (NONE), (ANYONE), SysAdminG, and AuditG, FieldServiceG, and OperatorG.

Reserved Group	Description
(NONE)	A group name used to specify that no one has access to an object except SysAdmin and members of SysAdminG. The group name is in all caps and includes the parentheses.
(ANYONE)	A group name that you use to specify that everyone has access to the object. The group name is in all caps and includes the parentheses.
SysAdminG	The name of the group whose members can read, write, and append/execute all documents, annotations, folders, queues, procedures, and so on, on the system. This group is a member of all groups.
AuditG	The name of the group whose members can read, but not write or append/execute, all documents, annotations, folders, queues, procedures, and so on, on the system.
FieldServiceG	A group created only for compatibility with previous releases. FieldServiceG has no special permissions.
OperatorG	A group created only for compatibility with previous releases. OperatorG has no special permissions.

Refer to "Group Assignments" on page 203 for the meaning of (NONE) and (ANYONE) when assigned as administrative, primary, or session groups.

Members of the SysAdminG and AuditG groups have no power to create and activate users, assign them to groups, modify their attributes, or change passwords unless the members also have the corresponding administrative attributes. See "User Security and Administrative Attributes" on page 207

Group Assignments

You can make users, devices, and groups members of an unlimited number of groups. In addition, you can make three special types of group assignments: administrative, primary, and session.

Administrative Group Assignments

The system assigns every security user, device, and group to an administrative group. An object's administrative group determines which security administrators can perform administrative tasks for the object (for example, activating an expired account). An object's administrative group also determines which administrators can perform administrative tasks for the object's extended membership and data object access.

A security object's administrative group is inherited from the administrator who created the object. Only SysAdmin can delete an administrator or change the administrative group of a security object. If the administrative group is (NONE), only members of SysAdminG can administer the object. If the administrative group is (ANYONE), anyone with administrative attributes can administer the object. Although the administrative group is assigned to the object, the object does not become a member of the administrative group.

You can set up your security system with the same groups as membership groups and administrative groups. However, you could create certain groups that are used only as administrative groups.

Primary Group Assignments

Every user is assigned one primary group and becomes a member of that primary group. The primary group assigned to a user determines who has access to various data objects created by that user. The data objects affected by a user's primary group include the following:

- annotations, notes, and highlights created through your desktop application
- folders
- WorkFlo systems, procedures, and queues
- forms and signatures on forms

Note

The document class security controls access to user-generated documents and batches, rather than the user's primary group.

The default primary group for users is (NONE). If you leave (NONE) as the primary group for a user, then no one except SysAdmin and members of SysAdminG has access to data objects the user creates. You can change a user's primary group to (ANYONE). All folders, annotations, and WorkFlo gueues then created by the user are accessible to everyone. Permissions for existing data objects created by a user are not changed when the user's primary group changes.

If the user's primary group is changed, the user retains membership in the original primary group, while acquiring membership in the new primary group. The administrator who assigns the primary group can

assign any group in the administrator's administrative domain (see "Administrative Domain" on page 197). The primary group for a group is itself, by definition, and the group's primary group assignment cannot be changed.

Session Group Assignments

The administrator assigns every user to a session group (groups and devices do not have a session group). The purpose of the session group is to permit an administrator to control logon privileges, such as logon times and expirations, for a group of users. By assigning a particular session group to several users, the administrator can add or modify logon privileges for the entire session group at once.

A session group assignment does not affect a user whose system attributes are set to override the system defaults, since the user values always override the session group values. The default session group is (NONE), which is equivalent to not having a session group. (ANYONE) cannot be assigned as a session group. The session group is assigned to the user, but the user does **not** become a member of the session group.

qiT To prevent users from logging on when you need to perform a task

such as defining indexes (see "Define an Index" on page 119), put all the users in a session group and expire the session.

Group Deletions

Before deleting a group, you need to be aware of what kind of data access or logon restrictions the group is providing to users. For example, if you delete the group that has read access to many of the documents on your system, many users would no longer have access to the documents they need.

If you delete a group functioning as a session group, then users that formerly had logon restrictions might suddenly have no restrictions or no access, depending on their override capabilities.

When you delete an object's primary group, that object's primary group reverts to (NONE), but membership in the group is not revoked.

Before deleting a group, you can review its members and see what groups it is a member of. However, no report shows you which objects depend on the group as a session or primary group.

You should not delete an administrative group without first using the Re-Assign Administrative Groups function. If you are not SysAdmin, then you must ask SysAdmin to run the Re-Assign Administrative Groups function. See "Reassign Objects to a Different Group" on page 248 for details.

If SysAdmin deletes an administrative group, the group attribute becomes (NONE) for all objects to which this administrative group had been assigned.

Deleting an administrative group has the following effects on the administrators for whom it was their administrative group:

- The principal administrator can still create objects but no administrator can see any of the objects.
- Only SysAdmin can administer objects with the group attribute (NONE).

Reports (summary, detail, extended membership) are blank. Administrators can still view current logons and event logs.

Note

You cannot delete the reserved groups: SysAdminG, AuditG, FieldServiceG, and OperatorG.

User Security and Administrative Attributes

If you are SysAdmin (or if the Allow System Override system default is checked), you can allow a particular user to override system defaults. For example, that user could set logon times which are different from the system default.

Many user permissions are granted through direct and extended group membership. A user is a member of the following groups:

- Primary group
- All groups of which you explicitly make the user a member
- All groups of which the explicitly assigned groups are members through extended membership

Always be aware of the extended membership of any group you make a user a member of.

Note

Assigning an administrative or session group to a user does not give the user membership within that group.

Classes of Users

The two broad classes of users are:

- Administrative users
- Nonadministrative users

Administrative users are distinguished by the assignment of one or more administrative attributes: supervisor, principal, group, and password. You can assign one or more administrative attributes to a user.

Four attributes give a user administrative abilities:

Administrative Attributes

Attribute	Description
Supervisor	Can update and delete security objects and log off users. Can manage users, groups, and devices belonging to the supervisor's extended group membership. Only an administrator with the supervisor attribute can view event logs and activate user accounts by turning off the expired status.
Principal	Can add security objects and log off users. When a principal creates a security object, that object automatically acquires the administrative group of the principal administrator. This group determines which administrator can manipulate the object. When a principal administrator without the supervisor attribute creates a new security object, the object is set to an expired status. A principal administrator without the group attribute cannot set or change a user's primary group.
Group	Can add members to and delete members from groups and functions, as long as the members are in the administrator's extended group membership.
Password	Can add or update passwords of non-administrative users within the administrator's extended group membership. Unless the password administrator also has the supervisor attribute, the user must be expired by a supervisor before the password administrator can change the password.

Some changes can be made only by administrators with a combination of attributes or by the SysAdmin user. SysAdmin is a special user with all administrative privileges and membership in all groups. Attributes are assigned on the Add User and Update User dialog boxes.

Only SysAdmin can create, modify, and delete administrative users. Administrative users cannot modify each other's characteristics.

In a small installation, one or two administrators might perform all administrative functions. The SysAdmin user can also perform these functions. In a large installation, one administrator might be in charge of keeping passwords and would need only the password attribute. Another administrator might be in charge of creating and updating users and would need principal and supervisor attributes.

Since group membership determines what security objects users can access, you could divide responsibilities for group memberships among several administrators who understand the details of their groups and can assign people to the appropriate groups. Administrative attributes do not control access to data objects, such as documents, folders, annotations, or queues. Only group memberships allow access to data objects.

Extensible User Authentication

Extensible User Authentication expands the capability of the IS security subsystem by enabling you to create a custom authentication library that is unique to your IS system. With Extensible User Authentication, the IS security subsystem uses the custom library to validate user credentials (user name and password) for your site, using an external authentication service instead of the Image Services MKF security database.

The external authentication service is typically a Lightweight Directory Access Protocol (LDAP) server, but it can be any third party authentication or Single Sign On product.

Extensible user authentication affects all security logons to Image Services with the exception of the following user categories:

Predefined IS Users:

SysAdmin, Operator, and FieldService If a conflict exists between a predefined IS user and an LDAP account, you must change the LDAP user name.

Unified Logon Users:

Unified logon users cannot log on locally to IS servers or use remote administrative tools like the Remote Admin Console (RAC) because the passwords of unified logon users were automatically generated when they were imported to IS.

Fallback Authentication

When the Fallback Authentication option is enabled, a valid IS user can log on to a local IS system even if the custom library rejects the credentials. In this case, the IS security service automatically performs a standard security check and validates the user's credentials against the IS MKF security database. The Fallback Authentication option is only available when you enable Extensible User Authentication.

Document Security

For each data object (document, folder, annotation, and so on) you can assign the following access privileges:

- User or group name that has read access
- User or group name that has write access
- User or group name that has append/execute access

The meanings of these access types depend on the data object (see table "Permissions" on page 254). You control access to data objects by putting the appropriate users in the groups that you assign to each type of access.

Document Classes

Batches and documents acquire security attributes from their document class. A newly created document class uses (ANYONE) for each type of access (read, write, append/execute). To secure documents in this class, change the security attribute to an appropriate user or group name.

Specify the security attributes for the documents in a document class when you create the class. If you change security for a class after committing documents to it, documents already committed have the old security attributes which differ from documents committed after the change. You can change security attributes for committed documents using the Update Document Security option in Database Maintenance (see "Update Document Security" on page 185). Each annotation, margin note, and highlight has its own security attributes. See "FileNet Notes" on page 212. In addition, each tab has its own security attributes. See "Tabs" on page 212.

Uncommitted Batches

Batches acquire security attributes from their document class. Operators who scan, index, reassemble pages, commit batches, or delete uncommitted batches need read, write, and append/execute permissions.

Folders

When a user creates a folder, the system assigns permission to read, write, and append/execute to the user's primary group. To see or perform any operation on the folder, a user must belong to that group. You can change the group. You cannot set append/execute for folders on a PC workstation. Your desktop application automatically uses the group with write access as the group name for append/execute.

FileNet Notes

Each FileNet note (which includes annotations, margin notes, and highlights) has three security attributes: read, write, and append/ execute access rights. When a user creates a FileNet note, permission to read, write, and append/execute is assigned to the user's primary group. The system checks note access rights when the user accesses notes in a display dialog box.

Tabs

Each FolderView tab has three security attributes: read, write, and append/execute access rights. When a user creates a tab, permission to read, write, and append/execute is assigned to the user's primary group. The system checks tab access rights when the user toggles between notes and tabs in a display dialog box.

Note

When users without a primary group create folders, notes and tabs, those objects will have access rights set to the user, not the group. If that user should be deleted from the security database, no other users will be able to access those objects.

Function Security

One of the system defaults is **Allow access to undefined functions** (see <u>"Set Up System Defaults" on page 228</u>). This allows you to secure only a few functions and leave all others accessible. For example, you might have several administrators that have access to the console, but you don't want them to access Database Maintenance. You can secure Database Maintenance by adding the **dbmaint** function code and assigning a group of which the administrators are not members. When administrators who are not members of the assigned group display the Application Executive's Applications menu, the Database Maintenance option is not visible.

If the **Allow access to undefined functions** box in the Update Default Security Settings dialog box is not checked, you must add each function to make it available. Only functions you define appear on menus, and they appear only to the members you specify. SysAdmin can perform undefined functions (not available from the menu) through the command line.

Function Security Planning (RAC systems only)

This feature is available with Remote Admin Console (RAC) systems only. For more information about RAC, see the *Remote Admin Console User's Guide*. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

The RAC interface translates all existing function codes into function descriptive names or roles. This is done through a "code-to-role" map-

ping mechanism. All function code displays have been replaced by role descriptions. There are two levels of role descriptions in Security Administration: application level and feature (function) level. Application level roles are designed for monitoring application accesses while feature level roles are tasks that can be performed for a particular application. This is called responsibility-based access control (RBAC). Listed below are the full relationships between the two levels (see also "Appendix A – Function Codes" on page 586)

Application Level Role	Feature Level Role
Database Maintenance	Archiving Utilities
	Delete Doc./Folder
	Update Doc. Security
	Update Retention Parameters
	Define/Update Index
	Rename Index
	Build Retrieval Key
	Drop Retrieval Key
	Define/Update Cluster
	Index Report
	Build Menu
	Define/Update Class
	Class Report
	Define/Update Family
	Family Report
Security Administration	NONE

Application Level Role	Feature Level Role
Storage Library Control	Detailed Surface Info
	Pending Surface Requests
	Media Space Usage
	Local/Foreign IDs
	Create Doc Header File
	Enable/Disable Storage Library
	Change Media Type
	Media Family Info
	Change Family Name
	Local Statistics
	Remote Committals
	Respond to RSVP
	Delete Info Message
	Configure Library screen
	Insert Media
	Eject Media
	Preformat Media
	Slot Drive Map
	Media Surface Summary
	Calibrate Library
	Identify Media in Library
	Media Family Information
	Eject Media by Location
	Delete RSVP
	MSAR Backup

Application Level Role	Feature Level Role
Background Job Control	Incorporate Foreign Media
	Manually Incorporate Foreign Media
	Copy Documents
	Copy Documents Using File
	Copy Annotations From Database to Media
	Consolidate Media
	Erase Media
	Rebuild Media
	Import Documents From Media
	Create Archive Database
	Import Archive Image Service
	Find Open Documents
	Completed Jobs
	Results of Find Open Documents
	Modify Status of Background Job
	Delete Log of Completed Job
	MSAR Convert
	Migrate Documents

Application Level Role	Feature Level Role
Cache Export/Import	Export Cache Objects
	Show Cache Objects
	Import Cache objects
	Cache Backup Program
COLD Main Menu	Define Background Template
	Define Channel Control File
	Define Report Format
	Define Import Job
	Preview Documents
	Import Documents
	View Import Log

Application Level Role versus Feature Level Role

This feature is available with Remote Admin Console (RAC) systems only. For more information about RAC, see the Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

It is important to remember that the purpose of application level roles is to monitor feature level roles. Think of an application level role as a gate. When the gate is closed (application not launchable), none of its associated feature level roles can even be accessed. When the gate is open, all feature level roles (menu options) behind that particular gate are now accessible and further security function checking could be imposed.

In general, RBAC should be achieved by activating application level roles only. However, in order to be fully effective, all feature level roles must also be activated. See "Case Examples" on page 224 and "Best Practice" on page 221.

Built-in Security Attributes

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

The 4 major security attributes (SUPERVISOR: Supervisor, PRIN-CIPAL:Principal, GROUP:Group, or PASSWORD:Password) within the Security Administration application are built-in features. They are not controlled by any external function codes or names. They are entirely controlled within the Security Administration application itself.

What Is a Role?

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC see the Remote Admin Console User's Guide.

A role is defined as a function access that performs a specific task(s). Roles are basically derived from "Function Codes." There are direct mappings of function codes to roles (see "Appendix A – Function Codes" on page 586). There is a role associated with each function code. The text describing what the role does is called the "Function Name".

For example, dbmaint is a function code that is used for controlling access to Database Maintenance application.

Function Code:	dbmaint
Function Name or Role:	Database Maintenance

A function code is always associated with a role or function name. A role is always associated with a function code.

Relationships Between Function Codes and Roles

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see the Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

The introduction of RAC provides a high level, user-friendly interface. Each role that a user activates enables one or more tasks for a particular application. In most cases, each role can perform only one task. However, there is one special occasion where a role (function code) allows users to perform multiple tasks:

Role	Allowable Task
Enable/Disable Storage Library	Enable/Disable Media
	Enable/Disable Library
	Enable/Disable Slot
	Enable/Disable Grippers
	Enable/Disable Drive

Default Behavior

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see the *Remote Admin Console User's Guide*. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

The default behavior on an Image Services install is that if a role hasn't been activated, then anyone has access to the role. This is determined by the global setting Allow Access to Undefined Functions in the Default Security Settings window. When a role is activated, everyone who should have access to it needs to be specifically added to that role (either directly or by group reference). Therefore, when a function is activated for the first time, anyone the administrators forget to add to the function's access list will suddenly lose access to the controlled application and its features. A suggestion is to initially put [ANYONE] in the controlled list until the administrators are comfortable that all authorized groups and users have been accounted for.

Allow Access to Undefined Functions

The global setting **Allow Access to Undefined Functions** is on by default. When this setting is checked (on) users are allowed access to all inactivated roles. If a role is activated, but with no members, then no one has access to that role. To allow a user access to an activated role, the user must be a member of that role.

When the global setting **Allow Access to Undefined Functions** is unchecked (OFF), no one is allowed access to inactivated roles. Granting access to a user is done by making the user a member of the activated role.

Special user: SysAdmin

SysAdmin is a special user who has access to all roles regardless of the state of the global setting.

Best Practice

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see the Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

In general, the best practice is to always apply security against groups or roles. Add users to one or more groups as appropriate. Only in special cases should security be granted directly to a user.

New System Considerations

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see the *Remote Admin Console* User's Guide.

It is very important for System Administrators to understand the usage, power, scope, and features of RAC. Please read "Default Behavior" on page 220 before the actual implementation takes place. It is suggested that you create a sample scenario where a test user and group are given certain responsibilities and see how the software behaves.

If a new system is being set up, it is easy to suggest what some common roles might be (for example, data center operator, application developer, help desk, application user, and application manager). In general, groups should be established to support roles that exist in the company. Based on the needs of each group, function access should be activated to support the desired level of security and appropriate

groups added to each function. New installations should activate each of the major application function names and [ANYONE] should be put in the function group if it is okay for anyone to run the program. Each user should get their own user ID and each user ID should be placed into the appropriate group or groups.

Existing System Considerations

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see the *Remote Admin Console User's Guide*. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

If the system already has function codes activated, all previously entered function codes now appear as role descriptions. The most notable GUI change is the **Activate Function Name** window, which is displayed via the Function pulldown. Users are no longer required to enter an accurate function code from "Appendix A – Function Codes" on page 586 of the System Administrator's Handbook. Instead, they will be asked to choose from a list of roles (see "Appendix A – Function Codes" on page 586). If users had previously entered inaccurate function codes, these invalid function codes will be appended with the string Unsupported Function on the console display. Unsupported function codes will show up in the Deactivate Function... window and the View Functions and Members..." window. The best thing for System Administrators to do is to delete these invalid codes.

RBAC is backward compatible to function codes. Although it is strongly recommend that users adapt to RBAC for its ease and user-friendly interface, users who prefer to see function codes as they did previously can still do so by enabling the Toggle Function Code Display option via the Function pulldown.

An administrator with an existing system must be much more careful, depending on what aspects of the system he wishes to control. Often on pre-existing systems, users are allowed access to any undefined function codes and probably have no function names activated. This means everyone can run anything assuming they can access a server. Once they activate a function, only those people who are in groups added to the function (or who are added to the function directly) will be able to run the application associated with the function name. Unless they already have a good mapping of roles to groups, some initial difficulty will result in resolving which groups and or individuals should be granted access. It may be the case that defining new groups and then adding existing users and groups to it would be the cleanest way to address this.

General Considerations

This section applies to Remote Admin Console (RAC) systems only. For more information about RAC, see the *Remote Admin Console User's Guide*. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

There are two primary cases that require that function access names be defined for all administrative applications:

- Access to only Security Administration by help desk personnel.
- Access to only Storage Library Control related functions by data center operators.

Any application without an activated function name is open for use by anyone. So, for example, to exclude the help desk personnel from everything except Security Administration means that function names need to be activated for everything. The help desk role-based group

would be added only to the **Security Administration** function name. The same applies to the data center operator group, except that their group is added only to the Storage Library Control function name (and probably also to the **Background Job Control** function name.)

In a day-to-day usage situation, administrators need help with adding new users and groups successfully. It is best to handle security control through groups, so choose functions that a group has access to when the group is created. Although this could be done for users also, it wouldn't really support "best practices" and could potentially encourage the administrator to adopt lazy security design habits.

Case Examples

See the Sample Scenarios chapter of the Remote Admin Console User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Device and Terminal Security

To use device security for printers and fax servers, your service representative enables security while configuring the device through the System Configuration Tools (see the System Configuration Tools online help). You need to know which names have been assigned to each printer or fax server.

Terminals, printers, and fax servers are members of groups and must have a common group membership. However, printers and fax servers are not subject to time restrictions and cannot be expired.

When you add a terminal, the system prompts you for a string name and the TCP/IP address of the device. The purpose of the address is to provide a unique name for the terminal. The address you specify here is not used for networking.

If you enable terminal security at the system default level, you must assign each workstation to a group appropriate for the users who need to use the workstation. If you do not enable terminal security at the system default level, then all devices are freely available to anyone unless you override the system defaults and turn on terminal security at another level.

If terminal security is enabled, a user's extended group membership is compared to the terminal's extended group membership to determine if the user has permission to log on at that terminal.

Password Controls


You can require that passwords be a minimum length (0 to 8 characters) or contain one or more nonalphabetic characters. You can also set them to expire after a configurable number of days (with an optional grace period with warnings to the user). Passwords are always encrypted within the security database.

Users can change their own passwords with the Change User's Password option (see "Set User Passwords" on page 272).

If you change the minimum password length in the system defaults (see "Default Security Settings" on page 231) to greater than zero, the next time a user changes a password, the No Password checkbox is disabled and the new password must match the new requirements.

Logon Process

The following flowchart illustrates the logon process.



Start the Security Administration Application

Choose Security Administration from the Application Executive's Applications menu. A new system is initialized with three users. These users and the initial passwords are listed below:

User	Password	
SysAdmin	SysAdmin	
Operator	Operator	
FieldService	FieldService	

A new system is defined with the following groups:

SysAdminG

OperatorG

FieldServiceG

AuditG

You can assign administrative users four types of administrative attributes (see "User Security and Administrative Attributes" on page 207), either separately or in combination.

After you create new users with different administrative attributes, you may also want to explore the menus and dialog boxes while logged on with one attribute at a time.

Throughout the program, you receive various messages in dialog boxes. When you cancel an operation, the dialog box often presents several buttons. You must click one of these buttons to indicate what you want to do next. In general, select Continue to continue with the operation or Cancel to return to what you were doing before the dialog box displayed.

Set Up FileNet System Security

This section describes the steps and functions for setting up a new FileNet security system (FileNet security is not the same as your operating system security). If you already have a system set up, the information in this section will help you modify your system for your particular needs.

Note

You must log on as SysAdmin to see all menu options and perform all functions.

After details on setting up and overriding system defaults, this section provides a brief discussion on how to plan your security setup, then continues with the processes required to set up the security.


Set Up System Defaults

Your site has security requirements that should be enforced across the entire system, with possible exceptions. Also, all or most groups, users, and devices need to have certain basic attributes. You can set up custom defaults at each of these levels.

Through the Update Default Security Settings dialog box, you can:


- set system-wide terminal security
- set system-wide function security
- set system-wide logon controls
- set system-wide password controls
- define system-wide security logging

Select Default Security Settings from the System menu to display the 1 Update Default Security Settings dialog box.



2 Click on checkboxes in the upper half of the dialog box to toggle options on or off.

3 In the lower half of the dialog box, enter values in the entry fields.



4 When you've finished changing your system-wide defaults, click the OK button to accept the defaults and close the dialog box.

Click the Cancel button to close the dialog box without making any changes.

Each setting on the Update Default Security Settings dialog box is described in the table that follows.

Default Security Settings

Setting	Description
Terminal Security	If the checkbox is selected, a user can access a terminal if the terminal and the user are members (or extended members) of the same group and time restrictions are met. If the checkbox is not selected (the default), the above security is not imposed unless the terminal, user, or session group has the override attribute selected and terminal security is checked in the System Attributes dialog box.
Allow Access to	If the checkbox is selected (the default):
Undefined Function	 Access to a function is granted if the function does not exist in the Security database.
	 Access to a function is verified if the function does exist in the Security database.
	If the checkbox is not selected, all functions must be defined and verified.
Password Special Character	If the checkbox is not selected (the default), passwords are not checked for special characters. If the checkbox is selected, a special character (other than A–Z or a–z) is required in each user's password and the minimum length of a password is 1 character. When set, restrictions apply to new or modified passwords. If the password restrictions change, existing passwords are not affected.
Log Successful Logons	If the checkbox is selected (the default), successful logons are logged in the security log file. Logging of successful logons can also be controlled at the individual user, session group or terminal level.
Log Failed Logons	If the checkbox is selected (the default), failed logons are logged in the security log file. Logging of failed logons can also be controlled at the individual user, session group, or terminal level.
Log Security Changes	If the checkbox is checked (the default), the events log describes a security change (add object, modify object, etc.) and shows the complete record before and after the change. Logging of security changes can also be controlled at the individual user, session group or terminal level.

Default Security Settings, Continued

Setting	Description
Allow Override of System Defaults	If the checkbox is not selected (the default), only SysAdmin can grant users, groups, or devices the permissions to override system defaults. If the checkbox is selected, administrators other than SysAdmin can define users, groups, or devices to override system defaults. The defaults that can be overridden display in the System Attributes dialog box for users, groups, and devices.
CFS-IS No Annotation Security Mapping	If selected, the per annotation security permission will not be federated to CE. The CFS-IS annotation will be federated to CE with read/write/AX permission of "(ANYONE)". This means all authenticated users on CE will have full control to the annotation, which includes the capability to delete the annotation from both CE and IS.
	If deselected, the per annotation security permission will be federated to CE. You must run SEC_map to map IS users/groups to CE Distinguished Names prior to enabling this option. IS must be restarted for the new setting to take effect.
Password Change Upon Reset	When a password is reset for administrative purposes, this option will force all users to change their password when they first log on. For more information, see "Mandatory Password Change After Reset" on page 273.
Enable Custom Password Validation	An external shared library can be used to verify that passwords are secure enough for the specific environment. This check box must be enabled before the shared library will be loaded. See "Custom Password Validation" on page 274. If you change the Enable Custom Validation setting, you must restart the IS software before the change will take effect. For more information about the shared library formats and information about creating your own external shared library, contact your service representative.

Default Security Settings, Continued

Setting	Description
Enable Extensible Authentication	An external shared library can be used to verify that user credentials (user ID and password) are secure enough for the specific environment. This check box must be enabled before the shared library will be loaded. See "Extensible User Authentication" on page 277. If you change the Enable Extensible Authentication setting, you must restart the IS software before the change will take effect. For more information about the shared library formats and information about creating your own external shared library, contact your service representative.
Enable Fallback Authentication	When this option is enabled, an IS user can still log on to a local IS system if the custom library rejects his credentials (name and password). The security service then performs a standard logon automatically and verifies the user's credentials against the MKF security database.
Password Minimum Length	Use this field to specify a minimum length for all user passwords. If Password Minimum Length is 0 (the default), a user may choose not to use a password. If it is set to a number between 1 and 8, then passwords are required and must be at least that many characters long. The maximum password length is 8.
	If you also check Password Special Character, then the Password Minimum Length field requires at least one special character. When set, restrictions apply to new or modified passwords. If the password restrictions change, existing passwords are not affected.
Password Renewal Days	Must contain 0 to 365 days. If the value is greater than 0 (the default), users must change their passwords periodically according to the interval set. For example, set this value to 30 days and users must change their passwords every 30 days.
Password Grace Period	This field is available only if the Password Renewal Days is a value greater than 0. The grace period designates the number of days before the Password Renewal date during which the user is warned when logging on regarding the approaching renewal time. If the password is not changed by the renewal time, the user's account automatically expires when the user attempts to log on. If an account is expired, SysAdmin or an administrator with Supervisor privileges must reinstate the account.

Default Security Settings, Continued

Setting	Description
Password Failed Attempts	Must contain a number from 0 to 100. Sets the number of times a user can fail to log on within the specified time period before the account expires. The default is 0, which allows an unlimited number of failed logon attempts.
Account Deactivated If Failed Attempt Exceeds <i>n</i> Minutes	This field is available only if you set the Password Failed Attempts to a value greater than 0. Must contain a 0 to 100 minute designation. If the number of failed logon attempts (set in Password Failed Attempts) occurs within the number of minutes specified, the account is expired. SysAdmin or an administrator with supervisor privileges must reinstate the account. For example, if the Password Failed Attempts field contains 3, and this field is set to 5 minutes, you are allowed 3 failures within 5 minutes.
Maximum Concurrent Logons Per User	Must contain a number from 1 (the default) to 10,000. Controls the number of concurrent logon sessions that a user can have at one time on one system. When the maximum is reached, an error is returned for subsequent logon attempts.
Logon button	Click the Logon button to display the Logon Times dialog box. The Logon Times dialog box allows you to set logon time restrictions for security objects. See "Override Logon Times" on page 244.

Set Up the Default Group Template

The default group template is defined at the system level and is the default for each group you add. Updating this template does not affect existing groups. Selecting Default Group Template from the System menu displays this dialog box:

X Update Defaul	X Update Default Group Template		
Group Name:	Group Name:		
Comment:	FileNet Group class default settings		
More Attribut	es: System Expiration		

The following table describes the options in the Update Default Group Template dialog box.

Update Default Group Template options

Option	Description
Group Name field	Contains the group name in this format:
	object:domain:organization
	You cannot edit this field.
Comments field	Contains up to 79 characters of user-specified text.
	Default text is:
	FileNet Group class default settings.

Option	Description
System button	Click the System button to display the System Attributes dialog box.
	See "Override Security Object Defaults" on page 245.
Expiration button	Click the Expiration button to display the Expiration Date dialog box.
	See "Override Account Expiration Date" on page 246.

Set Up the Default User Template

The default user template serves as the default for each new user created. Updating this template does not affect existing users. If you do not check the Override System Defaults/Session Group checkbox in the System Attributes dialog box, the system or session group defaults are used and system defaults in the template remain disabled. When you've finished changing the system-wide defaults, click the OK button to accept the defaults and close the dialog box. Click the Cancel button to close the dialog box without making any changes.

Selecting Default User Template from the System menu displays a dialog box similar to the following example.

X Update Default User Template		
User Name:	DserDefaults:moorea:FileNet	
Comment:	FileNet User class default settings	
More Attribu	tes: System Expiration	

The following table describes the options in the Update Default User Template dialog box.

Update Default User Template options

Option	Description
User Name field	Contains the user name in this format:
	object:domain:organization
	You cannot edit this field.
Comment field	Contains up to 79 characters of user-specified text.
	Default text is:
	FileNet User class default settings.

Update Default User Template option

Option	Description
System button	Click the System button to display the System Attributes dialog box.
	See "Override Security Object Defaults" on page 245.
Expiration button	Click the Expiration button to display the Expiration Date dialog box.
	See "Override Account Expiration Date" on page 246.

Set Up the Default Device Template

The default device template provides the default for each device you add. Updating this template does not affect existing devices. If system default overrides are allowed, the system attributes you set in the template override the system attributes set in the Default Security Settings. You can override these defaults when you create a new device. When you've finished changing your system-wide defaults, click **OK** to accept the defaults and close the dialog box. Click Cancel to close the dialog box without making any changes.

Selecting Default Device Template displays this dialog box:

× Update Default Device Template	
Device Name:	DeviceDefaults:noorea:FileNet
Device Class	Terminal
Comment:	FileNet Device class default settings
More Attributes: System Expiration	

The following table describes the options in the Update Default Device Template dialog box.

Device Template options


Option	Description
Device Name field	Contains the device name in this format:
	object:domain:organization
	You cannot edit this field.
Device Class field	Select the device that your template applies to from the pulldown list: Terminal, Printer, or Fax.
Comment field	Contains up to 79 characters of user-specified text.
	Default text is:
	FileNet Device class default settings.

Device Template options, Continued

Option	Description
System button	Available only for terminals.
	Click the System button to display the System Attributes dialog box.
	See "Override Security Object Defaults" on page 245.
Expiration button	Available only for terminals.
	Click the Expiration button to see the Expiration Date dialog box.
	See "Override Account Expiration Date" on page 246.

Override System Security Defaults

You can set system default overrides at the user, session group, and terminal level. The following illustration and table illustrate which of these settings take precedence.



Security Overrides chart

In general, the user settings take precedence over the session group settings, and the session group settings take precedence over the system settings. If you check the override checkbox for the terminal, the system uses the most restrictive settings between the terminal and user. The one exception is security logging, which occurs if set for the user, session, or terminal.

Override settings affect the following:

- Terminal security
- Time use restrictions
- Security logging
- Maximum concurrent sessions (user only)

Security Administration resolves which of the settings to use by determining whether the user, group, or terminal has permission to override the system settings.

Security Overrides


System Default Override Set for User	System Default Override Set for Session Group	System Default Override Set for Terminal	Result
Υ	Υ	Υ	Use most restrictive combination of user and terminal settings.
Υ	Υ	N	Use user settings.
Υ	N	Υ	Use most restrictive combination of user and terminal settings.
Υ	N	N	Use user settings.
N	Υ	Υ	Use terminal settings.
N	Υ	N	Use session group settings.
N	N	Υ	Use terminal settings.
N	N	N	Use system default settings.

You use three dialog boxes to specify settings for security objects.

Dialog Box	Where to Access
Logon Times	Accessible through the Logon button in the More Attributes section of various dialogs.
	See "Override Logon Times" on page 244.
System Attributes	Accessible through the System button in the More Attributes section of various dialogs.
	See "Override Security Object Defaults" on page 245.
Expiration Date	Accessible through the Expiration button in the More Attributes section of various dialog boxes.
	See "Override Account Expiration Date" on page 246.

Override Logon Times

In the Update Default Security Settings dialog box (see Step 1 on page 229), you can set logon time restrictions by clicking the Logon button. You will also see the Logon button whenever you access the System Attributes dialog box (see "Override Security Object **Defaults" on page 245**). When you click the Logon button, the Logon Times dialog box displays.



By default, logon times are unrestricted. To set system-wide logon time restrictions, click the Yes radio button. Then click the down arrows to select a "From" and "To" day of the week, hour of the day, and minute of the hour to limit access to the system. To return to the default, click the No button.

For example, you could define the system-wide default to allow most objects to access the system on Monday through Friday, from hours 7:00 through 19:00 (7:00 AM to 7:00 PM). Once these are set, you can define certain individual users, groups, or terminals to override the system settings. For example, a backup group might access the system Monday through Saturday from 7:00 PM to 7:00 AM.

Both the time and date are evaluated when determining whether a user can log on. For example, a user set to log on from 22:00 to 2:00, Monday through Friday, is allowed to log on Saturday at 1:00 AM, but not Monday at 1:00 AM.

Override Security Object Defaults

When you define a security object, you can override the system defaults by clicking the System button in the More Attributes section of the dialog box. The System Attributes dialog box displays.

X System Attributes ☐ Override System Defaults/Session Group □ Log Successful Logons **⊏** Log Failed Logons □ Log Security Changes

To override any of these system defaults or the time restrictions available from the Logon button, check the Override System Defaults/Session Group checkbox at the top of this dialog box. This makes everything in the dialog box available. The Maximum Concurrent Logons applies only to users. See "Set Up System Defaults" on page 228 for a description of each option.

When you change a system default, it might affect existing objects. For example, if you change the system default for the maximum concurrent logons allowed, it affects an existing user unless the user overrides system defaults and has a different number of maximum concurrent logons set.

Override Account Expiration Date

The Expiration button is available in the More Attributes section of user and group definitions (both the templates and the individual records). Clicking this button displays the following dialog box.

ОК	Cancel
More Attributes:	Logon
□ Terminal Security	

X Expiration Date	
↑ Use Expiration	
♦ No Expiration	
♦ Expired	
Year(nnnn): Month:	Day:
Hours: 0 Minutes: 0 S	econds: 0
ОК	Cancel

Selecting the No Expiration radio button activates an account created by SysAdmin or an administrator with principal and supervisor privileges. Other new accounts are created with an Expired status and cannot be used.

To make an expired account active, an administrator with permission to activate the account must choose one of these two radio buttons:

Radio Button	Purpose
Use Expiration	Assign an exact time when the account expires.
No Expiration	Prevent the account from expiring at a preset time. (The account may expire for other reasons, such as failed logons or an unrenewed password.)

To set an expiration time, click the arrows and choose an option from the list. For the Year, enter all four digits (for example, 1996).


Reassign Objects to a Different Group

SysAdmin can move a number of objects (users, groups, devices) from one administrative group to another. For example, use this option when the objects' previous administrative group is deleted or when a section is moved from one department to another.

Note

To change a single object to a different administrative group, updating the object directly is faster than using this option. For example, to change one user's administrative group, select Update User from the User menu.


1 Select Re-assign Administrative Group from the System menu, to display the following dialog box.



- 2 Click on the listbox arrow to display the Administrative groups.
- 3 Select the Administrative group that contains the security objects whose administrative group you want to reassign and click OK.

The Re-assign Objects' Administrative Group dialog box displays.

This dialog box displays all objects (users, groups, devices) assigned to the group you selected in the previous dialog box.



- At the top of this dialog box, use the pulldown list to select the adminis-4 trative group to which you are going to reassign the objects.
- 5 Select the objects from the left listbox that you want to reassign.
 - Click on objects to select them, or click Select All to select everything in the list. To deselect an object, click it again. Click the Clear button to deselect all objects at once.
- 6 With the objects in the left listbox selected and the proper administrative group set, click the Assign button.
 - Those objects move from the left listbox to the right listbox.
- Reassign members of the original group to a different group. 7
 - Select the new administrative group from the Administrative Group listbox at the top of the dialog box. Then select the objects from the left listbox and click Assign.
 - To undo the reassignment, select the objects from the right listbox and click the Undo button. Selection of objects works the same as described for the left listbox.

Plan Security Details

Careful planning before you begin the group, user, and device security setups will help you avoid problems and backtracking.

You need to determine the following:

- Which users need access to the system?
- What kind of work will each user need to do?
- What permissions do the users need to do their work?

Once you've determined the users' needs, you can plan administrative groups. You need to create administrators with attributes that allow them to appropriately administer their respective groups.

Create a chart or a diagram to help you plan your FileNet security.

The following example should provide you with a good idea of how to set up your own FileNet security system.

Example Scenario

You need to set up security for the Credit Card Approval division of a large bank. Here is the work flow:

- Credit card applications are received by mail, then scanned, indexed, and committed to the system under the document class APPS.
- After committal, the applications are routed to an Approval Representative (AR) who gathers information, places documents in folders, and decides whether the applicant is approved or denied credit.
- If the application is questionable, the AR sends the application to a Fraud Analyst (FA) who makes a decision to deny the application or return it to the AR for approval.

In this scenario, we want to set up security so that:

- Scanners and indexers can do their jobs.
- The committal supervisor can commit documents.
- ARs and FAs can access the committed APPS documents.
- ARs and FAs can file and unfile documents in folders and delete folders.
- ARs and FAs can create notes and read each other's notes.
- The System Controller (SC) can change attributes of documents for the APPS class.

Based on the above criteria, you need to create users and groups and assign the appropriate users to groups and subgroups.

The following table describes the various permissions as they relate to documents, folders, FileNet notes, and tabs. Refer to this table when setting up your FileNet system security.

Permissions

Object	Read	Read, Write	Read, Append/Execute	Read, Write, Append/Execute
Document	Retrieve, display, print document	Change document attributes	Create notes on document	Delete documents; scan, index, commit, delete batch
Folder	List folders in query options, copy folders	Change attributes	File and unfile documents in folders	Delete a folder
Note	Display notes and their text	Modify text, change attributes		Delete a note
Tab	Display tabs and their text	Modify text, change attributes, delete tab		

The following chart describes the users, their jobs, and the permissions each needs to do his or her job (Read, Write, Append/eXecute).

Users, Jobs and Permissions

			Permissions	
Name	Job	Documents	Folders	Notes
Claudia N.	Scanner	R, W, AX		
Kathleen T.	Approval Representative (AR)	R	R, W, AX	R, W
John K.	Indexer	R, W, AX		
Shirley K.	Indexer	R, W, AX		
Rene M.	Scanner	R, W, AX		
Kathi G.	Fraud Analyst (FA)	R	R, W, AX	R, W
Diane S.	Fraud Analyst (FA)	R	R, W, AX	R, W
Susan S.	Approval Representative (AR)	R	R, W, AX	R, W
Tom M.	Fraud Analyst (FA)	R	R, W, AX	R, W
Peggy M.	Committal Supervisor	R, W, AX		
Melanie C.	Approval Representative (AR)	R	R, W, AX	R, W
Mike C.	System Controller	R, W		

Sample Setup


The most efficient way to set up a security system is from the top down. Using our example scenario, follow the steps below to create a sample security setup:

- 1 As System Administrator, name the first group GROUP1 and make it the group that all other users and groups are extended members of.
- 2 Set up the groups STORAGE, ADD, and APPS as SysAdminG groups.
- 3 Set up the groups NOTES/FOLDERS, SCAN/INDEX, and SYSTEM CONTROLLER and assign them each to themselves (not to SysAdminG).
- 4 Make Mike a member of the System Controller group and give him principal and supervisor attributes so he can make changes to the attributes of the document class APPS.
 - Mike is, by extended membership, a member of the STORAGE and GROUP1 groups.
- 5 Make Susan an assistant administrator with all four administrative attributes (supervisor, principal, group, and password).

- Make Susan the principal administrator of the NOTES/FOLDERS 6 group who sets up all users for that group.
 - Susan sets up the FAs and ARs as members of the NOTES/FOLD-ERS group.
 - Susan gives Kathleen and Diane both principal and group attributes.
 - Susan gives Kathi and Melanie the supervisor attribute.
- 7 Make Peggy an assistant administrator with all four administrative attributes (supervisor, principal, group, and password).
- 8 Make Peggy the principal administrator of the SCAN/INDEX group who sets up all users for that group.

Peggy makes all scanners and indexers members of the SCAN/INDEX group. FAs, ARs, scanners, and indexers are, by extended membership, members of APPS, STORAGE, ADD, and GROUP1.

The following diagram illustrates the example security setup, with groups, users, and administrative attributes labeled (see table "Users, Jobs and Permissions" on page 255).



Set Up Group Security

You create all groups, including groups assigned as primary, session, and administrative groups, in the same way. Before creating groups, review the group template settings.

Note

To avoid confusing group names with user names, adopt a naming convention, such as ending all group names with an uppercase G, to distinguish groups from users in reports and event logs.


You must have certain administrative attributes to perform administrative tasks affecting groups. You must have the Principal attribute to create groups, the Supervisor attribute to update groups, and the group attribute to add members to groups. Administrative attributes are discussed in detail in "User Security and Administrative Attributes" on page 207.

The System Administrator can perform any and all administrative tasks for all groups.


Add or Update Groups

To add or update a group, select Add Group or Update Group from the Security Administration window's Groups menu. You must have the principal attribute to add groups and supervisor attribute to update group memberships.

- Do one of the following to add or update a group. 1
 - To add a group, select Add Group from the Group menu. Enter the new group name in the Add Group dialog box, change attributes as necessary, and click OK. The Update Group dialog box appears.



- To update a group, perform the following steps.
 - Select Update Group from the Group menu to display the Specify Group Name dialog box




Enter the group name in the Name field or click Query to search for the group name. See "Using the Query Feature to Select a Group Name" on page 268.

Note

The Domain and Organization field information you specify will only be utilized if you enter a name and click OK. It is not passed on to the Query function.

Click OK. The Update Group dialog box appears.

2 Make your selections and entries in the Update Group dialog box that displays.



- To add your group to another group: 3
 - Click Add next to the Member of Groups list.
 - Enter the group name in the Name field or click Query to search for the group name. See "Using the Query Feature to Select a Group Name" on page 268.
 - Click OK. The group appears in the Member of Groups list.
- 4 To delete your group from another group:
 - Select the group in the Member of Groups list. а
 - Click Delete.
- 5 To add a member to the group:
 - Click Add next to the Group's Members list.
 - Enter the group name in the Name field or click Query to search for the group name. See "Using the Query Feature to Select a Group Name" on page 268.
 - Click OK. The group appears in the Member of Groups list.
- To delete a member from a group: 6
 - Select the group in the Group's Members list.
 - Click Delete.
- 7 Once you are finished:
 - Click OK to accept your changes and close the dialog box.
 - Click Save to save your changes and leave the dialog box open.

- Click Next to clear the changes from the dialog box and make more changes (be sure you save any changes you want to keep before clicking Next).
- Click Cancel to close the dialog box without making any changes.

The following table describes the options in the Add (or Update) Group dialog box.

Add (or Update) Group options

Option	Description
Group Name field	Displays the group name in object:domain:organization format.
Comment field	Contains up to 79 characters of user-specified text. The default is:
	FileNet Group class default settings.
System button	Click the System button to display the System Attributes dialog box (see "Override Security Object Defaults" on page 245).
Expiration button	Click the Expiration button to display the Expiration Date dialog box. See "Override Account Expiration Date" on page 246.
Administrative Group field	Displays the administrative group name of the administrative user who created the group. Only SysAdmin can view or change the administrative group assignment.
Primary Group field	Displays the name of the group you created or updated. You cannot edit this field.
Member of Groups listbox	List all the groups that the group you are adding/updating is an extended member of (see "Extended Membership" on page 197).
Add button	Click Add to display the Group Selection List. You can choose from this list or enter a search pattern. When you enter the search pattern, a group (in the Group Selection List) which matches the search pattern is automatically highlighted. Click OK when you've selected the group to add.
Delete button	Select the member to remove from the group and click Delete.

Add (or Update) Group options, Continued


Option	Description
Group's Members listbox	Lists the members of the group selected in the Member of Groups list.
Add button	Click Add to display a list of security objects. Select those you want to be members of the group you are creating or updating and click OK.
Delete button	Select the member to remove from the group and click Delete.

Delete Groups

The procedure you use to delete a group depends on the group's assignment: Administrative, Primary or Session. See "Group Deletions" on page 205 for more information.

To delete a Primary or Session group

Select Delete Group from the Group menu to display the Specify 1 Group Name dialog box.



2 Enter the group name in the Name field or click Query to search for the group name. See "Using the Query Feature to Select a Group" Name" on page 268.

Note

The Domain and Organization field information you specify will only be utilized if you enter a name and click OK. It is not passed on to the Query function.

- 3 Click OK. The Delete Group dialog box appears.
- 4 Click OK to delete the selected group.

Note

You cannot delete the reserved groups: SysAdminG, AuditG, FieldServiceG, and OperatorG.

To delete an Administrative group

- 1 Run Reassign Administrative Groups (see "Reassign Objects to a Different Group" on page 248) to reassign the objects belonging to this group to another administrative group.
- 2 Follow the directions for deleting a Primary or Session group.

Update Group Membership

Group administrators can add members to and delete members from groups they administer. They can also make a group a member of another group, delete it from a group, or add a selected group in their administrative domain to other groups. If you have both group and supervisor attributes, you can use Update Group Membership to perform both kinds of functions.

- 1 Select Update Group Membership from the Groups menu.
- 2 Enter the group name in the Name field or click Query to search for the group name. See "Using the Query Feature to Select a Group" Name" on page 268.
- 3 The Update Group Membership dialog box appears. You can now add or delete groups and members. For more information, see "Add or Update Groups" on page 260.


Rename Groups

An administrator with the supervisor attribute can rename groups in his administrative domain.

1 Select Rename Group from the Groups menu to display the Specify Group Name dialog box.



- 2 Enter the group name in the Name field or click Query to search for the group name. See "Using the Query Feature to Select a Group" Name" on page 268.
- 3 Click OK. The Rename Group dialog box appears.
- 4 Enter the new group name in the Rename Object dialog box and click OK.



Using the Query Feature to Select a Group Name

If you aren't sure of the group's name, click Query to search by name. The Query for User Name dialog box appears. You have two options for searching by name:

- Choose Select All and click Submit. Select a name from the Match. List and click OK.
- Choose By Name enter a number in the Number in Match List field. Enter the first one or more letters of a group name in the Name field. The group names are case sensitive. Click Submit.

The Match List pane displays the first series of names based on the number you specified. Select a group name from the Match List and click OK.

Set Up User Security

After you set up your system security defaults and create groups, you are ready to add users. Only users can log on and perform tasks. The system provides a user template, which you can modify. A new user created by a principal administrator has an expired status. An administrator with the supervisor attribute must activate the logon account by changing the expired status. If you create a new user without a password, an administrator with the password attribute must set the password for the user before the user can log on.

Add Users

Follow this procedure to add a new user.

- 1 Select Add User from the Users menu.
- 2 Enter a user name (up to 40 characters) in the Name field of the Add User dialog box.
- 3 Use the pulldown list to select the domain and organization, which are separated by colons and are limited to 20 characters each.
- Click OK to display the Add User dialog box. Change the default infor-4 mation as appropriate (see table on next page).

The new user inherits the administrative group attribute from the administrator.

Note

The attributes you can give a user depend on your administrative attributes. For example, if you have the principal attribute, you can only create the user and change the Comment.

- 5 Click OK to accept your changes and close the dialog box.
 - Click Save to save your changes and leave the dialog box open.
 - Click Next to clear the changes from the dialog box and make more changes (be sure you have saved any changes before clicking Next).
 - Click Cancel to close the dialog box without making any changes.

The following table describes the options on the Add User dialog box.

Add User Options

Option	Description
User Name	The user name in object:domain:organization format. This field is grayed because you entered it in the previous dialog box and cannot change it here.
Comment	Contains up to 79 characters of user-specified text. Default text is:
	"FileNet User class default settings."
System button	See "Override Security Object Defaults" on page 245.
Expiration button	See "Override Account Expiration Date" on page 246.
Password button	See "Set Up Terminal and Device Security" on page 295.
Supervisor	SysAdmin checks this box to assign Supervisor status. Grants permission to update security objects.

Add User Options, Continued

Option	Description
Principal	SysAdmin checks this box to assign Principal status. Allows the user to create, but not update, security objects.
Group	SysAdmin checks this box to assign Group status. Allows the user to make an object a member of a group.
Password	SysAdmin checks this box to assign Password status. Allows the user to change other users' passwords within the administrator's extended membership.
Administrative Group	The name of the administrative group of the administrator who created the user. Only SysAdmin sees this field.
Session Group	Lists groups when Group box is checked. When a Session Group other than (NONE) (the default) is assigned, logon controls of the session group apply to the user, unless the user can override system default privileges. The user is not made a member of the session group.
Primary Group	Lists groups when Group box is checked. The primary group determines the default security access and can be changed. The primary group determines who has access to data objects, such as folders or annotations (but not documents), created by the user. The user is made a member of the primary group. The default is (NONE).
Member of Groups	Lists all the groups this user is a member of.
Add button	Displays the Group Selection List, which contains the same selection as the list button for Primary Group. Select from the list or enter a search pattern. When you enter the search pattern, a group in the list which matches the search pattern is automatically selected.
Delete button	Deletes groups from the Member of Groups list. Select the groups to delete from this list and click the Delete button.

Set User Passwords

Selecting Change User's Password from the Users menu displays the Change User Password dialog box. This dialog box will be slightly different for password administrators than for other users.

For users who do not have password administrator privileges, the Change User Password dialog box will only allow them to select the "No Password" option or enter a new password.

For password administrators, the Change User Password dialog box will contain the additional option, "Password Never Expires". For more information about the "Password Never Expires" option, refer to "User Expiration Exclusion" on page 274.

For all users, when you enter a new password, you must enter it twice. If the second entry does not match the first, you receive an error message that the password was not changed because the verification failed.

Note

Windows Server users, when you make changes to a user password, you must then run Application Executive and refresh that user's password information for unified logon. See "Unified Logon for Image Services for Windows Server" on page 81.

Extensible Password Authentication

Extensible password authentication in IS provides the ability to enforce stringent password validation rules and create customized password validation rules.

 Mandatory Password Change - The default security setting "Password Change Upon Reset" forces a user to change his password before logging in. This is required after the System Administrator has reset the user's password.

- **User Expiration Exclusion** excludes specific users from the password expiration rules.
- Custom Password Validation provides the ability to enforce more stringent, customized password validation rules.

The customer is responsible for providing a shared library defining the custom password validation rules. The library must have a single entry point designed to enforce restrictions on user passwords.

Note

Your Image Services system does not provide the external shared library. For complete information about creating your own customized password validation library, contact your service representative.

Mandatory Password Change After Reset

After the System Administrator resets a users password, and if "Password Change Upon Reset" is checked for this user, the user must reset his password before the next log in. This applies to all clients and administrative tools. The check box "Password Change Upon Reset" is located in the Security Administration Default Security Settings dialog. For more information about default settings, see "Set Up System" **Defaults**" on page 228. The default value for this feature is false.

This system-wide setting can only be changed by system administrators.

Note

System Administrators are exempt from the requirement to change their password after reset.

User Expiration Exclusion

The Change User Password dialog is shared by both password administrators and users when changing a password. The dialog has two different configurations based upon the type of user.

Both SysAdmin and password administrators will see the "Password Never Expires" check box on the Change User Password dialog box. They can set this value for other System Administrators and users. If this box is checked, the System Administrator or user will be excluded from the password expiration rules. The default value for the "Password Never Expires" field is false.

Users who are not system administrators can access the Change User Password dialog box but will only have the option to change their password.

Note

When a password administrator sets the "Password Never Expires" option for a user, and the system wide "Password Change Upon Reset" flag is set, the user will not be required to change his password.

Custom Password Validation

For added security, you can create a custom password validation function that can check for any set of conditions your installation requires.

For example, you might want users to have passwords that contain at least one numeric or one non-alphanumeric character. This feature allows for any type of password check and can vary widely from customer to customer.

To enable customized password validation, you must create a shared library to validate the user's new password. The IS software will load this shared library, if it exists, and the feature is enabled.

The customized rules validate the password and return either an OK status or an INVALID status. The IS system performs the external password check on all new passwords and password updates. This check requires that the password be re-entered if it is declared INVALID by the external shared library.

Note

If you change the Enable Custom Validation setting, you must restart the IS software before the change will take effect.

Enable or Disable Custom Library Validation

The shared external shared library feature must be enabled from the default security settings dialog. The default value is false, which means the external library will not be loaded or used for password validation.

Shared Library Entry Point

After you have enabled Custom Password Validation, the Image Services software uses the following entry point to link to your custom shared library:

where:

username is a null terminated string. The maximum length is 40 bytes.

password is an unencrypted null terminated string. The maximum length is 8 bytes.

You must create your own custom library function to do specialized checks. For HP-UX and Solaris systems, the library must have the name <code>libSEC_ext_valid_pwd</code>. For Windows and AIX® systems, the library must have the name <code>SEC_ext_valid_pwd</code>. The file extension varies between platforms:

AIX	SEC_ext_valid_pwd (no extension)
HP-UX	libSEC_ext_valid_pwd.sl
Solaris	libSEC_ext_valid_pwd.so
Windows	SEC_ext_valid_pwd.dll

For custom password validation, a valid password returns:

An invalid password returns:

The IS software uses the action parameter to identify the type of password change. The valid values for the action parameter are:

```
EXTERNAL_AUTH_PASS_ADD 1
EXTERNAL_AUTH_PASS_UPDATE 2
EXTERNAL_AUTH_PASS_DELETE 3
```

Important

For security reasons, you must restart the IS software after enabling the external shared library feature or making changes to the shared library itself.

Note

The Image Services system does not provide the external shared library. For complete information about creating a customized password validation library, contact your service representative.

Extensible User Authentication

Another optional security feature is Extensible User Authentication. You can create a custom library that can check for any set of conditions your installation requires.

To enable extensible user authentication, you must create a shared library to validate the user's credentials. The IS software will load this shared library, if it exists, and the feature is enabled.

The customized rules validate the user and return either an OK status (a value of zero), or a WARNING or ERROR status (a non-zero value). An INVALID status might cause IS to disable extensible authentication and revert to standard authentication.

Note

If you change the Extensible User Authentication setting, you must restart the IS software before the change will take effect.

Enable or Disable Extensible User Authentication

The shared external shared library feature must be enabled from the default security settings dialog. The default value is false, which means the external library will not be loaded nor used for user authentication.

Shared Library Entry Points

After you have enabled Extensible User Authentication, the Image Services software uses the following entry points to link to your custom shared library:

```
int SEC_IS_ext_initialize (char *scheme_URI_p,
                           char *config p );
int SEC_IS_ext_authenticate (char *scheme_URI_p,
                             char *username p,
                             char *password p,
                             char **result string pp,
                             char *TTL seconds p );
void SEC IS ext free result string (char *result string p );
```

The **SEC_IS_ext_initialize** entry point takes the following parameters:

scheme URI p is a null terminated string. The maximum length is 256 bytes.

config p is a null terminated string. The maximum length is 256 bytes. (Reserved for future use.)

After loading the custom library, IS calls this entry point only once. In some environments, the URI (Uniform Resource Identifier) and configuration options might be required.

Software developers can populate the scheme URI p variable to pass to the SEC IS ext authenticate entry point. Or the developers can choose to ignore this call and simply return a status of OK (a value of zero).

The **SEC_IS_ext_authenticate** entry point takes the following parameters:

scheme_URI_p is a null terminated string. The maximum length is 256 bytes. The value is obtained from the initialization call.

username_p is a null terminated string.

password_p is a null terminated string. The software developer is responsible for encrypting the password.

result_string_pp is passed in as a null string and is allocated by the custom library. If the library returns a result string, IS will log the data.

TTL_seconds_p represents the number of seconds (Time To Live) for the given credentials. This value is currently ignored by IS.

A non-zero return value indicates a warning or error condition. When the custom library returns a non-zero value, IS expects to find additional explanatory information in the result_string_pp field to post in the IS event log.

The SEC_IS_ext_free_result_string entry point is used exclusively for IS security to free any memory allocated to result_string_pp by the custom library.

You must create your own custom library function to do specialized checking and verification. For HP-UX and Solaris systems, the library must have the name **libSEC_ext_auth**. For Windows and AIX sys-

tems, the library must have the name SEC ext auth. The file extension varies among the operating systems:

AIX	SEC_ext_auth (no extension)
HP-UX	libSEC_ext_auth.sl
Solaris	libSEC_ext_auth.so
Windows	SEC_ext_auth.dll

Fallback Authentication

Fallback Authentication can only be enabled if Extensible Authentication is also enabled. Fallback Authentication provides for standard user authentication with the MKF database if the extensible authentication cannot authenticate a user.

Update Users

To update an existing user, perform the following steps.

- 1 Select Update User from the Users menu to display the Specify User Name dialog box.
- 2 Enter the user name in the Name field or click Query to search for the user name. See "Using the Query Feature to Select a User Name" on page 284.

Note

The Domain and Organization field information you specify will only be used if you enter a name and click OK. It is not passed on to the Query function.

- 3 Click OK. The Update User dialog box appears.
- 4 Update user values and click OK.

Only a system administrator with the supervisor attribute has access to the Expiration button. System administrators with supervisor and principal rights can access the Session Group field. If system overrides are allowed, the System button is enabled.

Changing the Name of the fnsw User

The "Programmable Security Objects" feature provides the ability to change the standard FileNet software user name fnsw.

Important

The task of changing any security object should be done with great caution. Before you actually change the fnsw user name, analyze your system, and if you can avoid changing this security object, do so.

If you decide to change this security object, review the new tools available to work with programmable security objects: fn_pso_driver, fn pso podf admin and fn pso switch in the Image Services System Tools Reference Manual. To download this manual from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Delete Users

A system administrator with the supervisor attribute can delete users.

- 1 Select Delete User from the Users menu to display the Specify User Name dialog box.
- 2 Enter the user name in the Name field or click Query to search for the user name. See "Using the Query Feature to Select a User Name" on page 284.

Note

The Domain and Organization field information you specify will only be used if you enter a name and click OK. It is not passed on to the Query function.

3 Click OK to delete the user.

Update User Membership

A system administrator with the group administrative attribute can make a user a member of one or more groups. A system administrator can also specify a user's group membership in the Create and Update dialog boxes.

- 1 Select Update User Membership from the Users menu to display the Specify User Name dialog box.
- 2 Enter the user name in the Name field or click Query to search for the user name. See "Using the Query Feature to Select a User Name" on page 284.

Note

The Domain and Organization field information you specify will only be used if you enter a name and click OK. It is not passed on to the Query function.

- 3 Click OK. The Update User Membership dialog box appears. Depending on your attributes, certain fields may not display.
- 4 To add a user to a group:
 - Click Add to display the Specify User Name dialog box.
 - Enter the user name in the Name field or click Query to search for the user name. See "Using the Query Feature to Select a User Name" on page 284.

- Click OK.
- 5 To delete a user from a group:
 - Select the group in the Member of Groups list.
 - Click Delete.
 - Click OK.
- 6 When you are finished:
 - Click OK to accept your changes and close the dialog box.
 - Click Save to save your changes and leave the dialog box open.
 - Click Next to clear the changes from the dialog box and make more changes (be sure you save any changes you want to keep before clicking Next).
 - Click Cancel to close the dialog box without making any changes.


Rename Users

An administrator with the supervisor attribute can rename users.

- 1 Select Rename User from the Users menu to display the Specify User Name dialog box.
- 2 Enter the user name in the Name field or click Query to search for the device name. See "Using the Query Feature to Select a User Name" on page 284.

The Domain and Organization field information you specify will only be Note used if you enter a name and click OK. It is not passed on to the Query function.

- Click OK. The Rename User dialog box appears. 3
- Enter the new user name in the Rename User dialog box and click OK. 4



Using the Query Feature to Select a User Name

If you aren't sure of the user's name, click Query to search by name. The Query for User Name dialog box appears. You have two options for searching by name:

- Choose Select All and click Submit. Select a name from the Match List and click OK.
- Choose By Name, enter a number in the Number in Match List field and enter one or more characters in the Name field. These characters should match the name you are searching for. Click Submit. Select a device name from the Match List and click OK.

Set User's Database Logon

If you have a full-use database license, you can use the Database Logons window to map an RDBMS user to an Image Services user for using embedded SQL commands on a PC.


Note

An RDBMS user must already exist before you can map it to an Image Services user.

You can set up the database logon permissions in several ways:

- Set individual database logons (one-to-one)
- Set a single database logon for several users (many-to-one)
- Set a combination of one-to-one and many-to-one logons


Select Database Logons from the Users menu to display the Database Logons window.



Add Database Logon

If you are the FileNet system administrator or if you have the principal attribute (see the table "Administrative Attributes" on page 208), you can add database logons.

1 Select Add Logon from the Logons menu to display the Add Database Logons dialog box:



- 2 Enter a database logon name.
- Enter and verify a password. 3
- 4 Click OK to accept your entry or Cancel to close the dialog box without making any changes.

Update Database Logon Password

If you are the system administrator or if you have the password attribute (see the table "Administrative Attributes" on page 208), you can update the logon password.

- Select a user name to update from the Database Logons window. 1
- 2 Select Update Logon from the Logons menu to display the Update Database Logons dialog box.
- 3 Enter and verify the new password.
- 4 Click OK to accept your entry or Cancel to close the dialog box without making any changes.

Map Database Logon to Image Services Users

If you are the system administrator or have both the principal and supervisor attributes (see the table "Administrative Attributes" on page 208), you can map the database logons to Image Services users.

- Select a database logon from the Database Logons window. 1
- 2 Select Map Logon to Image Service Users from the Logons menu to display this window:

imes Map Database Logon to IS User(s) $-$		
<u>F</u> ile <u>H</u> elp		
Database Logon Name: Native Database Pseudo Id:	saran 102	
Mapped Logons:	Unmapped Logo	ns:
Mapped IDM Image Services	ChrisS:moorea:FileNET FieldService:moorea:FileNET FieldService:moorea:FileNE Operator:moorea:FileNE SysAdmin:moorea:FileN SysAdmin:sundae:FileN SysAdmin:verde:FileNE saran:moorea:FileNET	enet Et Iet Et Et
Unmap >>		<< Map
Find Logon Mapped Unmapped Logon:		

- 3 To find a particular user:
 - Check the Mapped or Unmapped radio button.
 - Enter the name in the Logon field.
 - Click the Find button and the list scrolls to the entered name.
- 4 To map a user to or unmap a user from a database logon:
 - Select a name from the Unmapped Image Service Users or Mapped Logons list.
 - b Click the Map button to move the name to the Mapped Logons list or click the Unmap button to move the name to the Unmapped Image Services Users list.
- 5 When you finish all your selections, click OK to accept the changes or click Cancel to close the window without making any changes.

Delete Database Logon

If you are the system administrator or if you have the supervisor attribute (see the table "Administrative Attributes" on page 208), you can delete a logon name.

- 1 Select a name to delete from the Database Logons window.
- 2 Select Delete Logon from the Logons menu.
- 3 Click Delete at the prompt to delete the name, or click Cancel to close the prompt without deleting the name.

Set Up Document Security

Document security is determined when you set up your document classes in the Database Maintenance application. See the discussion under "Document Security" on page 211 and "Create Document Classes" on page 154.

Set Up Function Security

Use function security to identify which functions (menu items, buttons) are to be secured, then assign members to those functions, making them available only to those users. Members can be user names or group names. Before you decide how to assign functions, check the tables listing the function codes (see "Appendix A – Function Codes" on page 586). Some applications make many functions available within the application.

You can assign functions in two ways:

- To define a function and make it available to everyone, assign the group (ANYONE).
- To define a function and make it available to a restricted set of users, make the users members of the groups you assign to the function.

Security Administration has built-in function access. Only administrators can access most functions, and what is accessible depends on the attributes of each administrator.

CAUTION

Background Job Control's Erase Media function permits erasing optical media even if the media contain open documents. You may prefer to restrict this potentially hazardous function to a select few


administrators. See "Erasing Media" on page 406 for more information.

Add Function Names

Only SysAdmin can add function names. To add a function name:

- 1 Select Add Function Name from the Functions menu.
 - A popup window prompts for the name of the function.
- 2 Consult the tables in "Appendix A – Function Codes" on page 586 to be sure you are spelling and capitalizing the function name correctly, then click OK.

A dialog box similar to the following appears.



Click the Add button to display a list of user and group names. 3

- 4 Select one or more names that you want to be able to access the function and click OK.
- 5 To save the function and close the dialog box, click OK.

To save and continue adding more functions, click Save and Next. To close the dialog box without saving, click Cancel.

Activate Function Name

Only the System Administrator can assign user or group-specific functions (or select Activate Function Name from the Functions menu).

Note

If you reached the Activate Function window from the Update User or Update Groups window, Steps 1 and 2 below apply. If you reached the Activate Function Window from the Functions pulldown in the main Security Administration window, all of the steps below apply.

To activate a function name:

- 1 Next to the Show field, you have two choices:
 - Select the Application Level radio button and go to Step 2.
 - Select the Function/Feature Level radio button which will then activate the Feature Details area.

This area allows you to select a specific function feature to assign to this Group or User. Selecting any one of these Feature Details provides you with a variety of function name options to select in Step 2.

2 Select the desired function name you want to secure in the Choose Function Name pulldown and click OK.

Note

Consult "Appendix A – Function Codes" on page 586 to be sure you are selecting the desired function name correctly. Function names are grouped by levels. Select the appropriate level to see the associated function names.

If you came from the Update User or Update Groups window, you will return there.

If you came from the Functions pulldown in the main Security Administration window, continue with Step 3 below.

- 3 Click Add to display a list of user and group names.
- 4 Select one or more names to give those users or members of the group access to the function.
- 5 Click OK to save the function and close the window; click Save and Next to save and continue adding more functions; click Cancel to close the window without saving.

Update Function Membership

An administrator needs the group attribute to update existing function membership. To update a function membership:

1 Select Update Function from the Functions menu.

This displays a list of function names.

2 Choose the function you want to update (see "Appendix A – Function Codes" on page 586), then click OK.

A dialog box similar to that shown for adding functions displays, showing the current members.

- 3 To add more members, click Add and choose one or more names from the list.
- 4 To delete members, select them from the Function's Members list and click Delete.
- 5 To save the update and close the dialog box, click OK.

To save and continue updating more functions, click Save and Next. To close the dialog box without saving, click Cancel.

View Functions and Members

Choose View Functions and Members to see a list of the functions defined and the members assigned to them. The complete list displays for any administrator.

Deactivating a Function

You must be SysAdmin to deactivate a function. To deactivate a function:

- Select Deactivate Function from the Functions menu and the list of 1 defined functions displays.
- 2 Select one or more functions you want to deactivate and click OK.
 - A popup dialog box asks you to confirm your selection.
- 3 Click OK to deactivate the selected functions or Cancel to exit the dialog box without saving.

Set Up Terminal and Device Security

A new system has three default logon names: SysAdmin, FieldService, and Operator. The passwords are SysAdmin, FieldService, and Operator, respectively.

You can place a group membership requirement on a printer or fax server. You must know the name the printer or fax server was assigned when configured through the System Configuration Editor. You cannot assign logon times or expirations on a printer or fax server. Until you turn on terminal security, everyone has access to all terminals. You can turn on terminal security as a system default or in templates and object records and then define terminals to control their use.

You can approach terminal security in one of two ways:

- Turn on terminal security. Each user logon will fail, but the terminal is added to the security database. Administrators can then set appropriate security for each terminal.
- Leave terminal security off. Identify each terminal, set the appropriate terminal security features for each terminal, and then turn on terminal security.


Administrators need the same kinds of administrative attributes to manage devices as described for users and groups. If you have trouble accessing devices, turn off terminal security. (You may have to turn it off for more than one security object.)

CAUTION

If you have fax servers, do not use terminal security as a system setting. Instead, use individual terminal security. Do not set terminal security for fax server associated PC terminals, fax user logins, or primary groups associated with those users or terminals. The fax server will experience login failure even if the membership has been defined correctly.

Add Devices

1 Select Add Device from the Devices menu to display the Device dialog box:



2 Enter a device name.

> You must use the printer or fax server name configured through the System Configuration Editor. For terminals, you can use any name.

- "PC" is the device name prefix of a Desktop terminal. "WS" is the prefix for other Image Services client terminals.
- 3 Select a device class of Terminal, Printer, or Fax.
- 4 Click the TCP/IP protocol button and enter the terminal's address in the field to the right (use the format in the dialog box).
- 5 Click OK.
 - A new dialog box displays the device name and type. You can now set system attributes and administrative group, and make the device a member of the appropriate groups.
- 6 Click OK in this dialog box to save the device and close the dialog box; click Save and Next to save this device and add another; click Cancel to close the dialog box without saving.

Update Devices


To update an existing device:

- 1 Select Update Device from the Device menu to display the Specify Device Name dialog box.
- 2 Enter the device name in the Name field or click Query to search for the device name. See "Using the Query Feature to Select a Device" Name" on page 302.

Note

The Domain and Organization field information you specify will only be utilized if you enter a name and click OK. It is not passed on to the Query function.

3 Click OK. The Update Device dialog box displays.



Change the device attributes or group membership as described in 4 "Add Devices" on page 296 and click OK.

Note

You cannot update the device class or address. To change these values, you must delete the device and add it again.

Delete Devices

To delete a device from the security database:

- Select Delete Device from the Devices menu to display the Specify 1 Device Name dialog box.
- 2 Enter the device name in the Name field or click Query to search for the device name. See "Using the Query Feature to Select a Device" Name" on page 302.

Note

The Domain and Organization field information you specify will only be used if you enter a name and click OK. It is not passed on to the Query function.

Click OK. The Delete Device dialog box appears. Confirm to delete the 3 selected device.

Update Device Membership


To change the group membership of a device:

- Select Update Device Membership from the Device menu to display 1 the Specify Device Name dialog box.
- 2 Enter the device name in the Name field or click Query to search for the device name. See "Using the Query Feature to Select a Device" Name" on page 302.

Note

The Domain and Organization field information you specify will only be used if you enter a name and click OK. It is not passed on to the Query function.

3 Click OK. The Update Device Membership dialog box displays. Depending on your attributes, certain fields might not appear.



- Click Add to display the Specify Device Name dialog box. Enter the device name or use the Query feature described in "Using the Query Feature to Select a Device Name" on page 302. Click OK.
- 5 To delete the device from a group, click Delete to display the Specify Device Name dialog box. Enter the device name or use the Query feature described in Step 2.

Click OK. 6

Rename Devices

An administrator with the supervisor attribute can rename devices in his administrative domain.

- 1 Select Rename Devices from the Devices menu to display the Specify Device Name dialog box.
- 2 Enter the device name in the Name field or click Query to search for the device name. See "Using the Query Feature to Select a Device" Name" on page 302.

Note

The Domain and Organization field information you specify will only be used if you enter a name and click OK. It is not passed on to the Query function.

- 3 Click OK. The Rename Device dialog box appears.
- Enter the new device name in the Rename Device dialog box and click OK.

Using the Query Feature to Select a Device Name


If you aren't sure of the device name, click Query to search by name. The Query for Device Name dialog box appears. You have two options for searching by name:

- Choose Select All and click Submit. Select a name from the Match. List and click OK.
- Choose By Name, enter a number in the Number in Match List field and enter one or more characters in the Name field. These charac-

ters should match the name you are searching for. Click Submit. Select a device name from the Match List and click OK.

Task Flowchart

Use the following flowchart to help you determine whether or not the user can perform a task, based on the permissions owned by the user and, if terminal security is on, the terminal.



Security for Internetworking



To use resources of another system, you must log on to that system. Thus, you need a logon name that the other system recognizes. A logon name is a three-part NCH object. Usually you are not aware of using the domain and organization names, because the system supplies them automatically.

If you are logging on at system A to use resources on system B, you must log on to system A with a logon name that system B also recognizes. When you select a database or service on system B, the system tries to log you on to system B with the same three-part name and password you used to log on to system A. Thus, both systems need to have a common logon name and password.

To make it easier for users on system A to use resources on system B, include the system A logon name as one of those on system B.

SysA SysB

Kathi:SysA:FileNet John:SysB:FileNet John:SysB:FileNet Kathi:SysA:FileNet

Kathi can log on to SysA in the normal manner (using only her first name). The same is true for John on SysB. Both operators can use resources on the other system since each system recognizes the logon from the other system.

Each system is identified by a three-part server process name. The server process name is a reserved system object that is used to restrict interdomain access. A password is associated with the server process name.

If the server process names and passwords of two systems are the same, the systems can communicate with each other. Systems need to communicate, for example, to perform cross-system committal or to allow users on one system to access a printer on another system. The default is to access networked systems with the following name:

ServiceProcess:System:System

If you log on as SysAdmin you can change the name or password to stop intersystem exchanges in both directions.

Server process name restrictions apply only to FileNet services. As long as systems are networked, information can be exchanged at the operating system level. Even when the systems have different server process names or passwords, users on other systems can still log in and run operating system commands and selected FileNet tools.


The FileNet software checks function security on your local system even if you are accessing resources of another system. Object security is checked on the system that owns the data object (document, folder, note, or tab). To access a data object on a remote system, you must belong to a group defined on the remote system that has read access to the object. You must also have the appropriate permissions to perform other operations such as deleting and editing data objects.

Change Server Process Name



If your system is networked to other systems, you can log on as SysAdmin and prevent those systems from using your resources by changing your server process name.

Choosing Change Server Process Name from the System menu displays the following dialog box.




To change the name from the default, shown here, type another name and click OK. Once you save this name, other systems will not be able to access your system's resources.

To make your system's resources available to other systems, reset the server process name either to that shown in the illustration or to a name common to the systems that must communicate.

Change Server Process Password



Just as you can change your server process name, you can log on as SysAdmin and prevent other systems from using your resources by changing your server process password. Select Change Server Process Password from the System pulldown menu to display the following dialog box.



The following table describes each element of the dialog box.

Field	Description
User	Contains the current server process name whose password you want to change. You cannot edit this field.
Minimum Password Length	The default reflects the value set in Default Security Settings (in the System menu). You cannot edit this field.

Field	Description
No Password checkbox	If you check this checkbox, the system disables the Enter New Password and Re-enter New Password fields.
Enter New Password	Enter the new password for your system.
Re-enter New Password	Re-enter your new password to confirm correct input.

Security Reports and Logs

Security Administration provides reports that summarize and detail your security configuration. You can print these reports, save them to a file, or append selections to a file. In addition, the FileNet system logs security events. You can determine what information to log and then read the event logs through the Security Administration application.

Security Reports

The Users, Groups, and Devices menus include the following reports that you can view, save to a file, or print:

Report Type	Description
Summary	A list of three-part security object names
Detail	A complete record showing all attributes of the security object
Extended Membership	The extended membership list of the security object

As an administrator with the supervisor, principal, or group attributes, you can view information about all users, groups, and devices. After you have selected a view, you can change the view by selecting a different one from the View menu.

You cannot print reports if your primary group is (NONE).

An administrator with only the password attribute, regardless of group membership, can view the Summary reports, but not the Detail or Extended Membership reports.


Logon Reports


Selecting View Logons from the Users menu displays users currently logged on to the system. You can save all or portions of the report to a file or print the report. In addition, you can sort the information by user, by location (endpoint), and by time. To terminate a logon, select it and choose Kill Logons from the Logons menu. You will see no confirmation prompt, so be sure you're selecting the appropriate logon names.

Event Logs

Security event logs are updated daily for a period of 28 days. On the 29th day, the system overwrites the oldest event log. The system backs up those logs when you back up the dataset in the fnsw directory. Information in the event logs depends on what information you choose to log at the system default level, in all templates, and for all objects.


To view security events, choose Security Event Logs from the System menu. This displays a dialog box, as shown in the following example.





Only SysAdmin or administrators with the supervisor attribute can view event logs. All supervisor administrators see the same log information; it is not filtered according to administrative group or administrative attributes.

The list of event logs appears in a scrollable list, with the current day's log at the end. Select the day you want to view, then click OK to display the event log, as shown in the following example.



From the File menu, you can either save selected information in a file or print it. When saving to a file, you can either overwrite an existing file or append the information to the end of an existing file.

Locate a Directory or File

To locate a directory or file to append or save to:

- 1 Select the directory from the Directories listbox; the directory is reflected in the Filter edit field.
- 2 Click the Filter button. The files contained in the selected directory are displayed in the Files menu box.
- 3 Highlight the file name to place in the Selection edit field.

Save or Append the Event Logs

To save or append a portion of the log file or the entire log file:

- 1 Select the text you want to save either by dragging through it or by choosing Select All from the File menu.
- 2 Choose Save or Append from the File menu.
 - A Directories/Files dialog box is displayed.
- 3 If you know the file name and directory you are saving or appending to, enter that in the Selection edit field.
- 4 If you are saving to an existing file, you must confirm that you want to overwrite the file.

Print the Event Logs

- 1 To print a portion of the log file or the entire log file, first select the portion of the text you want to print.
- 2 Choose Print Selection from the File menu.
- 3 Complete the Print File or Report dialog box. (For details, see "Print File or Report Dialog Box" on page 543.)

System Management

This chapter describes system management tasks you regularly perform. You may distribute the various system management duties among several people.

Before describing the tasks, this chapter provides an overview of the system management programs you use most frequently to perform these tasks—the FileNet Task Manager and System Monitor.

Security administration, backup, and operating system management are system administration responsibilities that are described elsewhere in this manual or in your System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

For detailed information about security administration, see **Chapter 3**, "Security Administration," on page 194.

If you use EBR to back up and restore your system, see the Enterprise Backup/Restore User's Guide. To download this guide from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Operating system information is available in the documentation you receive with your operating system. However, certain aspects of operating system management related to your duties as the FileNet system administrator are discussed in the System Administrator's Companion

manual for your platform. To download this document from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

System Management Functions

The system management tasks described in this chapter are typical of those performed at most sites. Your site may require you to do more or fewer tasks than those listed. The primary system management functions are:

- Monitor the Image Services system
- Monitor event logs
- Monitor the flow of work
- Monitor storage use
- Remove unnecessary media
- Back up data to tape and verify your backup tapes. See the "Backup" chapters of these documents:
 - System Administrator's Companion for UNIX
 - System Administrator's Companion for Windows Server
 - Enterprise Backup/Restore User's Guide

To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

- Monitor background jobs (see "Current Jobs" on page 372)
- Maintain databases (see Chapter 2, "Database Maintenance," on **page 85)**
- Administer system security (see Chapter 3, "Security Administration," on page 194)

You must perform many tasks daily, while you might perform others only weekly, monthly, or as needed. Use the checklist in the following chart as a guide to determining when you might perform specific system management tasks.

Daily	Weekly	Monthly	As Needed
☐ Check error logs	☐ Check space in	☐ Delete error logs	☐ Power on/off
☐ Check filesystem	index database	☐ Delete unneeded files	☐ Train operators
space	☐ Check space in	from fnsw directory	☐ Maintain databases
☐ Back up databases	permanent and transient databases	☐ Delete core files	☐ Perform security
☐ Check SLC for	☐ Remove full	☐ Create system backup	updates
media requests	transaction log	☐ Remove old storage	☐ Defragment queues
☐ Check print queues	media and store offsite	media	and Visual WorkFlo tables
☐ Run batch session	Offsite	☐ Back up fnsw directory	13.5.00
reports			☐ Track and correct problems
☐ Monitor cache			☐ Review performance
			reports

FileNet Task Manager

The FileNet Task Manager provides a graphical user interface for controlling and monitoring the FileNet software.

You use the Task Manager to:

- Start, stop, and restart (recycle) the FileNet software
- Determine the active FileNet processes and the status of each
- Prepare the FileNet software for a backup or restore operation
- Examine event logs
- Monitor FileNet network connections

To launch Task Manager, follow the steps for your Image Services platform:



On a UNIX platform, log in as an fnsw user and enter the following command:


Xtaskman &



On a Windows Server platform, follow these steps:

- 1 Click the Start menu and select the Programs menu.
- 2 Select the FileNet Image Services option.
- 3 Select the Server Applications option.
- 4 Select the Task Manager option.

After launching Task Manager as described above for your Image Services platform, the main window appears.



The Task Manager main window displays the connected server's name, the current FileNet software state, and any active processes. Choose Refresh Rate from the Options menu to display a dialog box where you can change the refresh rate interval.

Note

The refresh rate you specify will apply to the current session. If you close the Xtaskman application and then restart it, the refresh rate is reset to the default value.

Monitor Functions

You can display the contents of event logs or monitor Remote Procedure Call (RPC) activity by making selections from the Monitor menu. For information about viewing event logs, see "Event Logs" on page 348. For a description of the RPC Activity report, see "Monitor **Network Activity" on page 332.**

Status Information

The main portion of the Task Manager window displays status information for the FileNet system. The Server field indicates the computer name of the server you are connected to. To connect to a different computer, type the new name in the field and click on the Connect button. If you have multiple FileNet Image Services servers, you can monitor and control a remote Image Services server through the Task Manager by specifying a different server in the Server field. The name you specify is an internetwork protocol (IP) host name (usually the same as your Image Services domain name).

The ability to connect to a remote host may be disabled through the optional Task Manager Configuration file. Either outgoing 'Connect' requests or incoming requests from other, remote servers may be blocked. If outgoing 'Connect' requests are disabled, the Server text field will be graved out. For more information, see "Appendix D -Task Manager Configuration File" on page 624.

The Software State field displays the current state of the FileNet software for the specified system. Typical entries show that the software has been started or stopped since a certain date and time. When failing to connect to a different server, the Task Manager displays a message in the Software State field.

Your service representative may ask you to determine the processes that are currently active on your FileNet system. The Current Processes table lists those processes along with information about each. The information in this table varies with the type of operating system. See "Check the Active Processes" on page 330 for more information about the contents of this table.

Software Control Buttons

To perform software control functions, click on any of the five control buttons at the bottom of the Task Manager window.

Click on the Start or Stop button to start or stop the FileNet software. Messages display in the Current Processes table and, at the completion of the task, the Software State field reflects the start or stop state.

Use the Restart button in situations where the prescribed action to solve a problem is to recycle the FileNet software (that is, to stop and then restart the software).

You can select the Backup Mode or Restore Mode button only when the FileNet software is stopped. Check the Software State field before you select Backup or Restore.

Selecting Backup or Restore shuts down the databases and initiates a minimal environment to run the backup or restore program. See the "Backup" and "Restore" chapters of these documents:

- System Administrator's Companion for UNIX
- System Administrator's Companion for Windows Server
- Enterprise Backup/Restore User's Guide

To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

System Monitor

The System Monitor window displays read-only reports about the state of the system. The reports are generated from data in the FileNet Management Information Base (MIB), the central Image Services status information facility. The System Monitor reports include:


- General system status information
- General user security status information
- Storage use
- Network activity
- Document services activity

The System Monitor automatically redisplays report information at intervals appropriate for the type of information displayed. To update the display in a report at any time, click on the Refresh button at the bottom of the report. You can save each report to a file. Select the Save As option from the File menu to display a standard file save dialog.

To print a report, click the Print button or select the Print option from the File menu. A standard print dialog displays where you select the printer, paper size, and other variables. (For details, see "Print File or Report Dialog Box" on page 543.)

Task Manager must be running before you can start the Application Executive (see "Starting Image Services Software" on page 79). Then, you can access the System Monitor through the Application Executive.

From the Application Executive's Applications menu, choose System Monitor. The System Monitor window displays.



The Storage menu has options for displaying magnetic disk cache, database, and storage library information. The Statistics menu provides statistics for networking connections and Document Services functions. For details on using these menus, see "View Storage Use Information" on page 338 and "View System Statistics" on page 344

The following tables describe the contents of the System Monitor window.

System Information

Field	Contents
Domain	Name of a FileNet system. See "NCH and Three-Part Names" on page 43 for details.
Organization	Name of an organization using a FileNet system. The name is FileNet unless your company has registered a different name with Xerox Corporation.
System Serial Number (SSN)	Your Image Services system serial number (not the machine ID)
Server Type	Type of server being reported. See "Image Services Servers" on page 38
System Up Time	Elapsed time since the FileNet software was last started

Security Information

Field	Contents
Users Currently Logged In	Number of users logged on to the FileNet system, including the console
Maximum Licensed Users	Depending on your system, up to 99,999 users can log on simultaneously.
Users Rejected Due to Exceeding License Limit	This field does not apply to the current software release.

Image Services

Table Column	Contents
Service	Each type of Image Services service running on the local system and used by client workstations. See <u>"Software</u> <u>Services" on page 45</u> for a description.
# processes	Number of server processes of a given type currently running
	The FileNet system starts a new process only if all currently running processes are busy.
max processes	FileNet system-defined maximum number of processes for each type
	When the maximum number of a type of process has been started, no more server processes of that type will be started, even if requests are pending for that service.
# rejects	Number of times a request for a new process of a given type has been rejected because the maximum number of processes of that type is running
	If this value is greater than zero, call your service representative for advice. You may need to increase the number of service processes defined for your system.

Monitoring the Image Services System

Initially, you should monitor your system to establish a baseline for media consumption rates and job processing throughput. Once you establish these baselines, continue monitoring at appropriate intervals, increasing the frequency as required. For example, monitor your media storage more frequently as your media becomes fuller, or check print queues more often if print jobs are not completing as quickly as usual.

You might need to collect and analyze system use reports. You can save to files many of the System Monitor reports that you run daily. You can collect these reports weekly to establish baseline activity and system use and to manage your system's resources.

To monitor the status of the software, you usually use Task Manager's Event Logs window (described under "Monitoring Event Logs" on page 327). Occasionally, your service representative might ask you to use other functions such as the Current Jobs table on the Background Job Control window (described under "Current Jobs" on page 372), the reporting capabilities of the System Monitor, or FileNet tools and utilities (described in Chapter 8, "Commands," on page 468).

Important

Verify that monitor programs are not scheduled to start during backup times. Starting IS processes while IS is trying to shut down to backup mode could cause the system to hang. If you create automated scripts or cron jobs that run IS processes, ensure that they are disabled before you shutdown IS for maintenence or backup.

Monitoring Event Logs

Events can be errors or normal activities like starting or stopping the FileNet software. When the FileNet system software detects an unexpected condition, it immediately generates a message. Some messages are merely informational and appear from time to time without cause for concern. Some messages indicate problems. Both types of messages are called events and are stored in event logs on the system.


The software generates an event log file each day. The oldest entry is at the beginning of the file and the most current entry is at the end of the file. Depending on your platform, the file name of an event log is el<yyyymmdd> or elog<yyyymmdd>. For example, the event log file for January 2, 1999 is el19990102 or elog19990102. This naming convention allows event log files to be sorted in date order.

At least once each day, you should check all event logs. Entries that contain words that indicate serious errors (such as failed, died, fatal, or crashed) and report these to your service representative. Also report event log files that are noticeably larger than usual.

Your system configuration can specify the location for event logs. By default, on each Image Services server platform, event log files are in the following directories:

Platform	Directory
Image Services for AIX/6000, HP-UX, and Solaris	/fnsw/local/logs/elogs
Image Services for Windows Server	\FNSW_LOC\logs\elogs

Select Event Logs from the FileNet Task Manager's Monitor menu to display the event log entries for the current day.



From the Event Logs' Display menu, you can choose one of two display modes: static or dynamic.

In static mode you can choose the current log or a log from a previous day. Click on the List button to list all the event logs, then select the

date of the log you wish to examine. Select another log with the Previous Log and Next Log buttons. Static mode displays new entries to the log file when you click on Refresh. Logs from other than the current day always display in static mode.

In dynamic mode, the system displays new entries as it adds them to the log file. Dynamic mode is valid only for viewing the current day's log activity and is useful for monitoring events as they occur. Click the Quit button (in dynamic mode, the Quit button replaces the Refresh button) to exit dynamic mode and reset the mode to static.

You can also use Task Manager to see event logs from a different server. In the Task Manager main window, change the server name to the one you want to view, then click the Connect button.

Gathering FileNet System Information

Occasionally, your service representative might ask you to provide information about your FileNet system environment. You might be asked to use functions of the FileNet Task Manager or System Monitor programs that you don't use regularly in order to help the service representative identify or correct a problem. In addition, to track or correct a problem, you could also be asked to collect data or monitor some aspect of your FileNet system.

In the following paragraphs, we discuss some of these functions, giving a brief description of reports and displays. You normally use these functions at the direction of a service representative who provides details about the procedure.

Check the Active Processes

You can see if certain processes are running on your system and you can display process information for currently active Image Services services. For these two activities, you use the FileNet Task Manager and the System Monitor respectively.

If the FileNet software is active, the Current Processes table in the FileNet Task Manager's main window displays the currently active processes and pertinent information about each.

The information in the Current Processes table consists of the name of the user who caused the process to be started, the process identifier (PID), a system-dependent identifier field, the process's start time, and the process name and its arguments.

🗙 FileNET Image Services - Task Manager File Options Monitor Help Server: Connect sundance Software State: Software started since 02/14/07 16:13:47

Current Processes:

User	PID	PPID	Start Time	Process
fnsw	1937	1	16:13:49	CSM daemon
fnsw	1931	1504	16:13:46	gti server
fnsw	1529	1504	16:13:17	/fnsw/bin/ilk_daemon
fnsw	1530	1504	16:13:18	MKF clean T
fnsw	1571	1504	16:13:18	MKF writer 0
fnsw	1932	1504	16:13:47	NCH daemon -pt
fnsw	2038	2037	16:13:50	OCOR Listen -pt - s32769 - t3600 - d20
fnsw	1933	1	16:13:48	SEC daemon
fnsw	1481	1	16:11:44	TIM daemon – s
fnsw	1934	1504	16:13:45	TM daemon ctl = f7 = p 0x5e0 = c 0x1
fnsw	1936	1504	16:13:45	TM daemon ctl = f7 = p 0x5e0 = c 0x1
fnsw	1939	1504	16:13:45	TM daemon ctl - f7 - p 0x5e0 - c 0x1
fnsw	1938	1504	16:13:45	TM daemon ctl = f7 = p 0x5e0 = c 0x1 TM daemon ctl = f7 = p 0x5e0 = c 0x1 TM daemon ctl = f7 = p 0x5e0 = c 0x1
				1
7	J			

The information in the system-dependent identifier field (between PID and Start Time) depends on your platform.

System-Dependent Identifier Field Contents


Operating System	System-Dependent Field Contents
AIX, HP-UX, Solaris	PPID (parent process identifier)
Windows Server	TID (process's thread identifier)

The information in the Current Processes table is refreshed at 30. second intervals. To change the interval, select Refresh Rate from the Task Manager's Options menu and change the number of seconds.

If additional process-related information is needed, the FileNet Services area of the System Monitor's main window displays the currently active Image Services services. See "System Monitor" on page 322 for details on the information displayed in the System Monitor's main window.

Monitor Network Activity

Your service representative may ask you to monitor network connection activity for the FileNet software. Select the RPC (Remote Procedure Call) Activity option from the Task Manager's Monitor menu. The server software must be running. If the server software has not been started when you select RPC Activity, a message reminds you to start the FileNet software.



By default, the system refreshes information in the RPC Activity window at two second intervals. To change the interval, select Refresh Rate from the Options menu and change the number of seconds. The system processes many RPCs in less than one second. These RPCs occur so quickly that they may not display in the RPC Activity window.

The following table briefly describes the contents of the columns of the RPC Activity report. Your service representative can provide further details about the data in these fields.

RPC Activity Field Descriptions

Column	Description
State	State of the network connection
Туре	Connection type (X in the column indicates a server connection; a blank column indicates a client connection)
PID	Process identifier of either the request handler (if this is a server connection) or a client program
COR PID	For server connections, this is the process identifier of the paired COR_Listen process handler linked to the request handler. For client connections, the value in this field is always zero.
Program	For server connections, this is the name of the request handler for the connection. For client connections, the name represents the request handler that the client is communicating with on the remote server.
RPC Time	Number of seconds since an RPC was transmitted or received. For servers, the time is reset each time an RPC is received. For clients, the time is reset each time an RPC is transmitted.
Address	Network address of the machine on which the RPC originated (server), or of the remote machine being contacted (client)

Monitoring the Flow of Work

You must monitor the flow of work through your system, making sure operators complete jobs and the system processes batches without delays. For example, you might need to move people from scanning to indexing if indexing cannot keep up with scanning. Otherwise, batch cache can fill, halting system operation.

As operators become more experienced, you can eliminate some optional tasks such as image or index verification. You should continue to monitor the progress of personnel and the entire document entry process.

Sessions

We suggest that you assign one person to the task of defining sessions (see "Session Group Assignments" on page 205). This gives you control over the naming conventions and the number of batches created at any given time. For more efficient system performance, do not define more than 25 sessions at a time. Review your plans with your service representative.

Design a session header sheet that best reflects your needs. The purpose of the header is to list each document entry task, the date it was completed, who completed it, number of documents and number of pages, and high and low Document IDs for each session. Session headers should remain with the documents at least until committal.

Batches

A hardcopy batch session log might help you locate a document in the future. Try keeping batch session log sheets in a binder on-site. The log might list session names, entry dates, page counts, and document ID ranges.

If more documents than expected accumulate in batch cache or page cache, the available media space can fill up, halting document entry. To prevent this and to prevent having to back up cache at the end of the day, check the status of batches at least once per day. Make sure that documents are being committed regularly and that no pending write requests exist (pending write requests remain in cache). You can find this information in batch session reports. See the "Document Entry" chapter of your desktop application's user's guide for information about batches and reports.

If uncommitted images are in cache at the end of the day, you should back up cache.

Use CSM exim (see "CSM exim" on page 474) or the Cache Export/ Import Program to back up locked objects such as documents only if you want to export your cache data to import on another Image Services server. Because this program does not synchronize the cache with the transient database, it isn't the best tool to use for backing up cache to restore for disaster recovery on the same machine.

If you want to back up your cache to store for disaster recovery on the same Image Services server, you should use the FileNet Enterprise Backup and Restore (EBR) program instead. EBR synchronizes cache with the transient database, thus ensuring the restored objects have the same IDs they had when you performed cache backup. (For details on using EBR, see the *Enterprise Backup/Restore User's Guide*.)

See the "Backup" chapter of the System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server for information about the Cache Import/Export Program.

To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Print Jobs

If you print or fax many documents, monitor the print queue to watch for problems that could cause delays in job processing and completion. For example, an out-of-paper condition can interrupt your print and fax jobs. You can display the print queue at a PC workstation. This information is not available from the console.

The information in print queue reports shows you the request ID, requester's login name, printer name, priority of the job, the time it was submitted, number of pages, the status of the request, and the type of job (printer or fax).

The PRI tool's requeststatus command displays slightly different information (see "requeststatus" on page 521): the request ID, total pages, number of pages printed, status of the job, and the printer name.

The printing subsystem deletes completed jobs from the queue after one minute and canceled jobs after 10 minutes. For more information about printing and viewing the queue reports, see the workstation user's guides and manuals that came with your print server.

Monitoring Storage Use

You should regularly monitor storage use to prevent running out of space and thereby halting system operation. The following discussion points out the areas you need to monitor, the tools you can use to perform these monitoring tasks, and some information about how to interpret the data you collect.

The FileNet system software should not vary much in terms of space use, but you need to monitor the system components that consume storage space and attend to any potential problems. Resolving problems may include adjusting the use of your storage space, clearing areas of storage, storing files to tape, or purchasing additional storage media.

A partial list of storage-consuming components includes:

- objects (including documents) stored in cache
- databases
- event logs
- temporary files
- core (dump or Dr. Watson) files

Magnetic disk can be consumed by documents that are not migrated to storage media. Cache gradually filling up may be normal (for example, when you keep all documents in cache and never migrate to storage media), but a full cache may indicate a problem. Since event logs accumulate daily and are not automatically deleted, they can consume your available space over time.

View Storage Use Information


To determine how your system uses storage in your system, you can examine the following System Monitor reports:

- Magnetic Disk Cache Information
- Database Storage Information
- Storage Library Information

Magnetic Disk Cache Information

Check cache space periodically throughout the day by using either CSM_tool or the System Monitor's Magnetic Disk Cache report. To use the CSM_tool utility, see "CSM_tool" on page 480.

Select Magnetic Disk Cache from the System Monitor's Storage menu to display the Magnetic Disk Cache report window:



Use the horizontal and vertical scroll bars to see all the data.


The Magnetic Disk Cache Information report displays the following data.

Local Caches

Column	Contents
Cache Name	Three-part NCH name of each logical cache
Min./Max. sectors	Minimum and maximum space allocation, in sectors, for each logical cache. When configuring cache, your service representative sets upper and lower limits to guarantee each individual cache partition some space, but prevent it from consuming the entire portion of disk allotted to cache.
Free sectors	Amount of unused cache space
Locked sectors	Number of sectors that cannot be reused until the information is processed. Generally this applies to objects (such as documents) that have not been written to storage media. An object can take up several sectors of magnetic disk space.
In Use sectors	Number of sectors that contain objects (locked and not locked)
Locked Objects	Number of locked objects in the cache. Documents not written to storage media are locked in page cache.
In Use Objects	Total number of objects that exist in cache (locked and not locked)
Description	Commonly used name for each cache

Database Storage Information

You can monitor individual database storage use by selecting Databases from the System Monitor Storage menu. The following window displays.



Use the horizontal and vertical scroll bars to see all the data. The Database Storage Information report displays the following information.

Note

The index database row will only appear on FileNet-controlled systems (FileNet software starts and stops the RDBMS database), not on sitecontrolled systems.

Local Databases

Column	Contents
Database	Name of each system database
Туре	Database type. All databases, except index, are Multi-Keyed File (MKF) type. The index database type you've chosen (Oracle, DB2 or MS SQL Server) appears here.


Local Databases, Continued

Column	Contents
Location	Full path to the location of each database on magnetic disk
Total (KB)	Amount of magnetic disk space allocated for the database, in kilobytes
In Use (KB)	Actual amount of magnetic disk space used, in kilobytes

Storage Library Information

Normally, you manage your storage libraries through the Storage Library Control (SLC) program, but you can monitor the status of your libraries and other important information from the System Monitor. For more information about SLC, see Chapter 7, "Storage Library Control," on page 419.

Check your library status by selecting Storage Libraries from the System Monitor's Storage menu. The system presents information about your storage libraries in the following window.



Use the horizontal and vertical scroll bars to see all the data.

The Storage Library Information report displays the following information.

Storage Libraries

Column	Contents
Library ID	Letter assigned to the library (a-z)
Library Type	Type of storage library

Storage Libraries, Continued

Lib Status	Status of the library:	
	Enabled Functioning normally	
	Disabled Library inoperable	
	Manual Library's mechanical arm disabled. However, you can manually insert and remove media from the drives.	
	Manual is the normal mode for an optical disk unit (ODU).	
Tot Drives	Total number of drives in the library	
Enabled Drives	Number of drives available. If a drive is disabled (using Storage Library Control), the number of enabled drives may be fewer than the total drives.	
Arm Moves	Total number of times the mechanical arm moved since the FileNet software was last started	
Media Loads	Number of times media were inserted into the library since the FileNet software was last started	
Media Unloads	Number of times media were removed from the library since the FileNet software was last started	

View System Statistics

Use the System Monitor's Networking Statistics and Document Services Statistics reports to collect information about your system. Typically, your service representative asks you to obtain this information to assist in troubleshooting, performance analysis, or similar tasks.

Networking Statistics

To view a report of network activity, select Networking Statistics from the System Monitor's Statistics menu. The screen displays statistics for events which occurred since the FileNet software was last started.

X FileNET Image Services - Network Statistics		
<u>F</u> ile <u>H</u> elp		
Incoming Courier Connections Approved: Incoming Courier Connections Rejected: Courier Connections Which Timed Out or Terminated Abnormally:		
Courier Connections Aborted:		
Outgoing Courier Connections:		
Failed Outgoing Courier Connections:		
*** Note: The Network Statistics are reset each time the IS software is started.		
Print Refresh		

The following table describes the networking statistics fields.


Networking Statistics

Field	Contents
Incoming Courier Connections Approved	Number of incoming connections to the server for all services
Incoming Courier	Number of incoming connections rejected for one of these reasons:
Connections Rejected	Image Services or operating system network failures
	An Image Services service received an invalid request
	System reached the maximum number of busy server processes of the required type. Errors of this type usually occur in the server, but could be caused by client or network error.
Courier Connections Which Timed Out or Terminated Abnormally	Number of connections that did not terminate in an orderly way. Abnormal termination could be caused by turning off a PC or other remote client without shutting down the PC's client applications.
Courier Connections Aborted	Number of connections aborted when the Image Services server encountered errors processing a request from a remote client (an RPC)
Outgoing Courier Connections	Number of connections a server opened to communicate with another server. The value in this field can also represent Image Services Toolkit (ISTK) applications running on the server.
Failed Outgoing Courier Connections	Number of connections a server failed to open to communicate with another server. A failed connection could indicate a communication problem on your server, trouble on the remote server, or a network problem.

Document Services Statistics

Document Services handles movement of data to and from storage media. From the System Monitor, you can view a report on pages and documents migrated to magnetic disk, magnetic disk and optical drive hits, committed documents and pages, import statistics, and fast batch committal statistics.

Select Document Services Statistics from the System Monitor's Statistics menu. The following window displays.



The statistics in the Document Services report are for events that happened since the FileNet software was last started, as described in the following table.

Document Services Statistics

Field	Contents
Pages Migrated to Magnetic Disk	Number of pages moved from storage media to magnetic disk to satisfy retrieval requests
Documents Migrated to Magnetic Disk	Number of documents for which one or more pages have been migrated to magnetic disk
Magnetic Disk Cache Hits	Number of times a request was satisfied by finding a document in cache
Optical Drive Hits	Number of times a request was satisfied by finding the document on storage media already in a drive
Pages Committed	Number of pages committed to the permanent database
Documents Committed	Number of documents committed to the permanent database
Documents Read from Imported Media	Number of documents read in from storage media during an import operation
Documents Imported to System	Number of documents committed to the permanent database by the import operation
Batches/Pages/Documents Committed Using Fast Batch Committal	Number of batches, pages, and documents committed using fast batch committal (see "Fast Batch Committal" on page 67)

Remove Unnecessary Media

Periodically, you should remove unnecessary storage media from the storage library. Unnecessary media may be those which you have consolidated onto other media or which will be moved to another site (for example, multiple copies of transaction logging media).

Manage Other Files

Event logs and other types of files can accumulate, consuming your available storage. Examples include temporary files that are never deleted and files created due to system or application program errors (known as core or dump files in UNIX-based systems and Dr. Watson files in Windows Server systems).

Event Logs

The system creates event logs daily. They contain entries for normal system activities and for error messages. Periodically backing up accumulated event logs to tape and deleting them from media is the safest management approach. You then have a tape copy for future troubleshooting.

For instructions on copying event logs to tape, see the "Backup" chapter of your System Administrator's Companion for UNIX, System Administrator's Companion for Windows Server, or the Enterprise Backup/Restore User's Guide.

To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

Core Files in UNIX-Based Systems



Do not allow core files to accumulate. Core files may have valuable information that can help solve system problems. Examine your system regularly for core files and copy them to tape (maintaining a copy is vital if the files are needed for troubleshooting). Then you can delete the core files on media, freeing up storage space.

Note

The system may dump a core file in any directory. When a core file exists, check the event log file for a message indicating the location of the core file, usually in /fnsw/local/tmp. A core file has the file name **core** or **AEcore**. Therefore, each core file in a directory overwrites any previous core file.

Dr. Watson Files in Windows Server Systems



Dr. Watson files in the Windows Server environment take up little media space. The system creates them during system failures. You can delete Dr. Watson files when you no longer need them for troubleshooting.

Database Server Connect

Database Server Connect is a utility that you can run through the Application Executive (Xapex) program. Login to Xapex > Applications > Database Server Connect. It is designed to manage the security issues (authentication)pertaining to the four required relational database users.

These are the four types of relational database users:

- **f sw:** The primary user of the Image Services relational databases and eProcess.
- f maint: Mainly used by GDB exim, a generic database export/ import utility. It is also used by your service representative to gain access to the system's relational database for troubleshooting and investigation.
- **f_sqi**: This legacy user is used by the SQI subsystem of IS Toolkit. If you have IS Toolkit or IS Process Analyzer installed, the f sqi user can be used to access some of the features of these products.
- f open: The default database logon user used by the SQI subsystem of IS Toolkit. It is the per user database logon default user.

Overview

It is an industry standard that operating system passwords must change periodically. Some companies periodically require changes to the database password. Changing passwords is just one component of a site's security policy. The Database Server Connect application is

designed to manage this policy for the relational database user passwords that are set during the initial configuration at the time of installation on the Root/Index and Application servers.

Image Services integration with the relational database requires the IS system to maintain the connection accounts and the associated passwords. These are used to enable IS to connect with the relational database.

For DB2 only

If the operating system password must change, both the operating system password and the IS maintained password for the specific account must change at exactly the same time.

CAUTION

Changing one password before the other will result in failed connections between IS and the relational database during the time interval when the passwords are different. In most IS installations, any time interval where this occurs is unacceptable.

To address this, the IS system maintains two passwords for each account, the Primary password and the Secondary password. During typical operation, you will use the Primary password for each account when connecting to the relational database. You use the Secondary password maintained by Image Services only when a connection with the account Primary password fails. In this case, you should try connecting to the relational database with the Secondary password. If the attempt succeeds, this Secondary password replaces the old Primary password and all future database connections use this new Primary password.

With this mechanism, changes in the operating system password will not result in failed connections. To accomplish this, set the IS maintained Secondary password to the new operating system password before any operating system password changes are made. After the Secondary password on the IS system is updated, the operating system password may be changed. The IS system will first attempt to use the old Primary password. This will fail because the operating system password has been changed. Automatically, the IS system will attempt connecting to the relational database using the Secondary password. This will succeed with the Primary password on the IS system being updated with the Secondary password value. It is best that you set a new value for the Secondary password by checking with the Database Administrator for the next remote database server password and then using the New Password window. See "Changing the Primary and/or Secondary Password" on page 355.

To change a database user password in DB2, you need to change it in two places. First, you need to change it on the operating system. Second, you need to change it on an Image Services through Database Server Connect.

For Oracle and SQL Server

You do not have to ask the DBA to change passwords for you. Database Server Connect will change passwords on both the relational database and Image Services for you.

For All Relational Databases

The relational database passwords (f_sw, f_maint, f_sqi, f_open) must be set up after the IS configuration is finished. In addition, when you change or otherwise manage the passwords, the Database Server

Connect application in Xapex also needs to be run on each IS server, Root/Index server and Application server because the passwords on each of these servers are independent.

You must be logged on with System Administrator level privileges in order to use the Database Server Connect application and change the relational database user passwords.

Database Connect Administration

The Database Server Connect window allows you to easily change passwords for the f_sw, f_maint, f_sqi and f_open users and keep track of the status of the Primary password with Activation and Expiration date information on each server where the relational database client is installed.

This window contains four tabs, each one corresponding to one of the database users. Each tab has a Primary area and in the case of DB2, a Secondary area.

Primary Area

- Password: field and Change button. The Primary password always has an assigned value. For instructions on changing the password, see "Changing the Primary and/or Secondary Password" on page 355.
- **Activation Date:** Displays the date the Primary password was most recently changed.
- **Expiration Date (DB2 only):** Displays how many days from the current date until the Primary password expires. This field is set to a system default of 60 days.

It is a good idea to periodically check the **Expiration Date** field as part of your password maintenance routine to know when it is time to change these passwords. For more information about password maintenance, refer to the For DB2 section in "Password Maintenance" on page 357

Note

For DB2, you will get a reminder message to change your f_sw Primary password. See "Expiration Notification (DB2 only)" on page 356 for more information.

Secondary Area (DB2 Only)

When the Primary password expires and a new connection is attempted, the Secondary password automatically becomes the Primary password, if it is the same as the current password set on the remote database server. The **Activation Date** field is automatically changed to the current date and there is no value in the Secondary password field.

Password: field and Change button. If there are asterisks displayed in the field, there is a Secondary password assigned. If the field is blank, there is no password and one should be assigned as soon as possible.

This Secondary password enables you to logon to the system after the Primary password has expired or if the remote database server password has been changed. If the Secondary password is not assigned and the Primary password expires, you will not be able to log in as that user and the Image Services software may not come up. See "Password Failure Emergency Procedures" on page 359.

Note

It is key to have the Secondary password set to what will be the next remote database server password.

Regardless of the user, the Primary Password is the current password. For DB2 only, the Secondary Password is the password that will take affect if the Primary Password is allowed to expire.

Change the password(s) by completing the following Steps:

- 1 From the main Xapex screen, select the **Database Server Connect** option from the **Applications** pulldown menu.
- 2 Select the tab for the user whose password you want to change.
- 3 Click the **Change** button for either the Primary or Secondary password to display the **New Password** dialog box.
- 4 In the **New Password** dialog box, enter a new password in the **New Password:** field. This password can be a string of any alpha-numeric characters.
- 5 In the Confirm Password: field, enter the same password that you entered in the previous step.
- 6 Click OK.
- 7 If you changed the Primary password, verify that the Activation Date has been changed to today's date.

Note

For Oracle and SQL the password is updated in both the relational database and Image Services. For DB2 the password is updated only in Image Services. The user is responsible for updating the password in the DB2 database.

Expiration Notification (DB2 only)

All Primary passwords have expiration dates, but you will only be reminded before the f sw password expires. The FileNet application Executive (Xapex) automatically notifies the System Administrator a certain number of days before the f sw user's Primary password expires. The default value is 14 days. The System Administrator will see a pop-up message stating "Your database 'f_sw' Account Password will expire in # days". This notification pop-up will appear each time a new logon to Xapex is made if it is within the value set in fn edit.

In fn edit, two values are set for all four database users (f sw, f maint, f sqi and f open) at the time Image Services is installed: **Password** Expiration Policy and Notify Administrator.

The **Password Expiration Policy** is a reminder for the users on the IS system that the password on the remote database server is going to change soon. This value is set to a default of 60 days, but should be consistent with the site's IT policy as it relates to the remote database server. Notify Administrator indicates when to start signaling the System Administrator at each logon sometime toward the end of the Password Expiration Policy. This value is set to a default of 14 days (mentioned above). Both of these values can be changed on the Root Index server in fn edit on the appropriate tab of the **Relational Data**bases tab. If you run the same instance of Xapex indefinitely without logging off and logging back on, you may never receive an expiration notification and the f sw password will expire. If this occurs and there is no set Secondary Password, refer to the "Password Failure Emer**gency Procedures" on page 359.** For more information about password maintenance, refer to the For DB2 section in "Password Maintenance" on page 357.

Note

When setting the Password Expiration Policy, a blank field is not permitted and a value of 0 is equivalent to **Never Expires**. This means that as far as Image Service is concerned the password on the remote database server is never going to change, so the user will never be prompted that the IS server's Primary password is going to expire. In this case, the value set in the Notify Administrator field is meaningless.

Password Maintenance

Note

For Oracle and SQL Server, there is no expiration notification mechanism implement in Image Services. Your System Administrator needs to remember when the database passwords are going to expire and change the passwords before they expire.

If your site has a password expiration policy for your relational database accounts, follow the procedures below to ensure that Image Service continues to work properly.

For Oracle and SQL Server

- Before the password expiration date, your IS System Administrator 1 must run the Database Server Connect application and set the new password for the account.
- 2 The expiration policy will automatically apply to the new password and Image Services will continue to run without any issues.

For DB2

It is important that proper database authentication for the IBM remote database server and its associated IS server(s) be maintained for security reasons and potential server accessibility issues.

We recommend performing the following password maintenance steps:

- 1 When setting the **Password Expiration Policy** values in fn edit during software installation, make sure you have an estimate of how often the passwords on the remote database server will be changed (per your site's IT policy). Armed with this information, you can set the expiration of the database user passwords on all of the IS servers that will be connecting to the remote database server to coincide with the changing of the remote database server's passwords.
- 2 When you receive the notification message that the f sw password is going to expire, make sure that Secondary passwords for all four f * users have been set to what the new passwords on the remote database server are going to be. This is should be done on all servers that are going to be accessing the remote database server.
- 3 Request that the Database Administrator change the four f * passwords on the remote database server to match what the Secondary passwords have been set to.
- 4 After the Database Administrator has changed the four passwords on the remote database server, the next time those users on the IS server(s) try to access the database server, the Primary password, which is about to expire, will no longer match, the Secondary password will automatically become the Primary password, and the log on will be

successful. This also resets the **Activation Date** to the current system date and the Password Expiration Policy cycle begins again.

Note

You should be aware of what the next password is going to be on the remote database server so you can set this as the Secondary password on all of the servers where the database client is installed, and that are connecting to the remote database server.

When no passwords succeed, refer to the "Password Failure Emer-5 gency Procedures" on page 359.

Password Failure Emergency Procedures

IS applications that are running before password expiration will run without issues. New applications or transactions that require a database connection will fail since Image Services is not aware of the password expiration and they will continue to try on logon to the relational database with the expired password.

Consequently, if your IS System Administrator does not set new password before the expiration date and the passwords expire, you will have to complete additional steps to sync up the new password between Image Services and the relational database as detailed below.

For Oracle

You must use Oracle sqlplus to change the Oracle password to a temporary password first. Then you must use the Database Server Connect application to set a new password.

- Use sqlplus to set a new temporary password for all four users (f sw, 1 f_maint, f_sqi, and f_open):
 - Start sqlplus:

sqlplus <user name> /<old password

- Oracle displays a message that the old password has expired and prompts you for a new password. Enter a temporary new password.
- 2 Invoke the Database Server Connect application.

Xdbconnect -r

- 3 When prompted, logon as the System Administrator.
- 4 In the Database Server Connect window, change the Primary password to match the password set on the remote database server for all four users as described in "Changing the Primary and/or Secondary Password" on page 355.
- 5 Restart the FileNet software to verify that the IS software successfully comes up.

For SQL Server

You must use the SQL Server tool **osql** to change the SQL Server password to a temporary password first. Then you must use the Database Server Connect application to set a new password.

Note

You can also use the SQL Server Management Studio to change the password.

- Use osql to set a new temporary password for all four users (f sw, f 1 maint, f_sqi, and f_open):
 - Start osql:

osql -s <server name>/<instance name -E

You should type the following at the prompts:

1> alter login f sw with password='<new temporary PW>';

2> **go**

- Repeat sub-step b for remaining users.
- Invoke the Database Server Connect application. 2

Xdbconnect -r

- 3 When prompted, logon as the System Administrator.
- 4 In the Database Server Connect window, change the Primary password to match the password set on the remote database server as described in "Changing the Primary and/or Secondary Password" on page 355.
- 5 Restart the FileNet software to verify that the IS software successfully comes up.

For DB2

- 1 Have you System Administrator change passwords on the remote database server for all four users (f_sw, f_maint, f_sqi, and f_open).
- 2 Run the **Xdbconnect** -r command.

- When prompted, logon as System Administrator. 3
- 4 In the Database Server Connect window, change the Primary password to match the password set on the remote database server as described in "Changing the Primary and/or Secondary Password" on page 355.
- 5 Set the Secondary Password to the next Database Administrator issued password for the remote database server, if known.
- 6 Restart the FileNet software to verify that the IS software successfully comes up.

Database Reconnect

Remote Oracle databases that use Real Application Cluster (RAC) can be stopped and restarted without having to restart the Image Services software. When the connection between IS and the remote database is lost, IS automatically attempts to reconnect.

Transactions that were in progress at the time of the lost connection will be automatically connected to a new session when the database becomes available. No data is lost.

IS attempts to reestablish the connection every five seconds for a total of 300 seconds (five minutes) before it gives up and logs a message in the elog. You can adjust these values by using the IS System Configuration Editor, fn edit.

Note

If you use the standard default values, these values do not appear in the fn edit display.

Modifying the database reconnect default values

1 To modify the default values or to turn off the database reconnect feature, start the IS System Configuration Editor:

fn edit &

- 2 On the Procedures tab, select Add Relational Database Object.
- 3 When you are prompted to **Enter the name that you want to use for** the RDB object, enter the name of one of the three available settings:
 - db reconnect disabled
 - db reconnect timeout
 - db reconnect interval
- 4 When you are prompted to **Enter the location of the RDB object**, enter the appropriate numeric value for the setting you want to modify:

db reconnect disabled 1 (disabled) or 0 (enabled) db reconnect timeout **nn** (default is 300 seconds) db reconnect interval **n** (default is 5 seconds)

To modify all three of these settings, you must run this procedure three Tip times.

When you finish, exit from fn edit and run fn build -a to put the new values into effect.

Background Job Control

Certain jobs in the FileNet system take considerable time to complete and are designed to run in the background. When a job is running in the background, you can still perform other administrative activities with the Image Services software. When a job is running in the foreground, the system is busy until the foreground job is complete. Background Job Control (BJC) is the interface you use to start, control, and monitor Image Services background jobs.

While you can request any number of jobs, they may not run immediately. For example, only one import job per server can be run. If you request a second import, the program queues it for later execution.

BJC queues any background jobs that cannot run because too few storage media drives are available. It may gueue any number of jobs, up to the limit of the MKF database. We recommend at least two drives to copy, consolidate, or rebuild media.

Note

If you disable a drive after a job starts running, BJC pause the job until you re-enable the drive.

Note

Background jobs such as Import, FindOpen Docs, and Doc Copy can cause excessive disk swapping, even though you set up fn edit parameters to minimize disk swaps. This is because background jobs circumvent those parameters. If you notice excessive disk swapping, check for background jobs. You can cancel and re-run them during a less active time.

Use Background Job Control functions to perform the following actions:

- Copy documents.
- Copy annotations from magnetic media to storage media.
- Incorporate foreign media into your storage library or ODU (not a background job, but included in BJC because incorporating foreign media is a step in importing documents from storage media).
- Import a specified set of documents from storage media.
- Find open documents on storage media in preparation for consolidating media.
- Archive records in the index and permanent databases to reduce database size and still keep the associated documents available for retrieval.
- Consolidate media.
- Erase media.
- Migrate documents when cache-only system adds a storage library.
- Convert optical surface to MSAR surface.
- Monitor current jobs.

Note

You only run Background Job Control on systems with storage libraries (optical disk or MSAR) or optical disk units (ODUs). Unless converting a cache-only "OSAR-less" system to optical or MSAR media, Background Job Control functions are not used by cache-only systems.

Number of Background Jobs

Because some background jobs are heavily database bound, the following rules exist:

- Document import jobs The maximum number of concurrent import jobs is one. If a second import job is requested, the job is queued for a later execution.
- Erase jobs Only one MSAR erase job and two optical erase jobs per server can be run at the same time; the rest will be gueued.
- Document copy jobs This is limited by the number of drives/2.
- MSAR conversion jobs This is limited by the number of drives.

Background Job Algorithms

All Jobs, except copy jobs, require only one drive. As long as there are drives available, the job will start. In general, the maximum number of concurrent jobs is the same as the number of available drives with one exception: Only two import jobs are allowed to run concurrently.

Copy jobs uses two drives per job. This is true for copy documents, consolidate, rebuild, and erase.

In general, the maximum number of concurrent jobs is the same as the number of available drives divided by two. There are, however, two exceptions:

Exception 1: If there is only one enabled drive and it's not being used by any background jobs, copy job will be started too. This means during the copy operation, the source and destination surfaces will be swapped in/out to/from the only drive if they are different.

Exception 2: If the number of available drives is the same as 1 + the number of drives being used by any background jobs, don't start the copy job until there are more drives freed up from other background jobs.

Note

Available drive refers to a drive that is configured and enabled, but not being used.

Background Job Files

Each background job you start receives a sequentially numbered job number and a log file. The log file name indicates the type of job (ImpLog for import, CpyLog for copy documents, AntLog for copy annotations, etc.). The last part of the name indicates the background job number. For example, ImpLog.000003 is an import of documents running as background job number 3. These ASCII files reside in the fnsw local directory. You view, print, and delete them using this program.

Checking Status of Jobs

Jobs that are in progress appear on the main BJC window. Once a job completes, use the options on the Completed menu to view the results.

Copying Media

Storage media (usually optical disks) contain image documents as well as additional information about alternate copies of documents (transaction logs), the serial number of the system that wrote the media, and surface numbers. You can duplicate the media, copy some of the documents, and copy certain information from magnetic media to storage media (annotations and more current index information).

Copying storage media enables you to maintain copies of documents at a separate location for disaster recovery or for an additional backup copy. You can also create backup copies of documents during committal by defining additional transaction logs. Each document is then written to the primary family and to each transaction log media. Or you can commit documents to a remote system as well as your own system.

Copying media is a two-part operation. The system reads from the source surface into the retrieval cache (up to 128 documents at a time, depending on cache space) at prefetch, or lowest, priority. Then, the system writes from cache to the new surface at normal system write priority, lower than retrievals, but higher than prefetches. This two-part operation means the system can copy from one surface to another even when only one storage media drive is available.

The copy operation also writes 60 bytes per document and 12 bytes for each error in the fnsw directory. Check your media free space before beginning a copy operation (see "View Storage Use Information" on page 338).

Copying does not create a mirror image of the source surface. If the destination surface has too many flaws to accommodate all documents on the source surface, the system writes to a second destination. None of the copy functions copies deleted documents.


If you start a document copy operation while the source surface has pending write requests, the pending writes will not be copied. Check to ensure the surface you are copying has no write requests against it. You can find the status of the write requests by viewing the Pending Surface Requests report from the Storage Library Control's Media menu (see "Pending Surface Reguests" on page 459).

A copy operation automatically restarts if the system or the FileNet software goes down during the operation. Since the copy operation restarts from its approximate position when the system went down, you may have duplicate documents on your destination surface. This duplication has no effect on the copy operation or the FileNet software.

Starting Background Job Control

To access Background Job Control, you must first log on to the FileNet software (see "Starting Image Services Software" on page 79 for procedures).

Once you are logged on, select Background Job Control from the Application Executive's Applications menu. The main window displays.



The main window consists of a menu bar at the top, Current Jobs table in the center, and control buttons at the bottom. These areas and their functions are described below. The display refreshes automatically at preset intervals.

Menu Bar

The Background Job Control window's menu bar contains five menus.

- The File menu provides ways to exit from the Background Job Control program and to print or save Current Jobs reports. (For details on using the print dialog box, see "Print File or Report Dialog Box" on page 543.)
- The Incorporate menu provides functions for incorporating foreign media into your system.
- The New menu provides functions for starting new background jobs.

Note

If you use any of these jobs to write data to an incompatible media surface (2X when using an 8X drive), the system rejects the command and displays an error message.

- The Completed menu provides functions for seeing the results of completed jobs.
- The Help menu provides access to online assistance for the Background Job Control functions.

Current Jobs

The Current Jobs table of the main window displays information about each currently running background job. Incorporating media (see "Importing Documents" on page 374) is a foreground job and is not shown in the Current Jobs table.

Not all columns are used for each type of job. If a column is not in use for a particular job, N/A (not applicable) displays in that column. The contents of each of the columns is described in the table below.

Column	Contents
Job #	Number the system assigns to the job
Туре	Kind of background job (copy documents, copy annotations, import documents, find open documents)
Status	Current status of the job (suspended or running)
Start Time	Time the job started
Phase	If a job uses phases, the number of the phase currently running
Errors	Number of errors encountered
% Done	Percent of the job that is complete
Current Objects	Number of documents that have been processed from the start of the job until the current time
Total Objects	Total number of documents to be processed

Control Buttons

The lower portion of the main window presents control buttons you can use to manipulate the background jobs shown in the Current Jobs display.

To perform an action, select one or more jobs from the Current Jobs (selected jobs become highlighted), then click on the button. Each button's function is described in the table below.

Button	Function
Suspend	Suspends the currently selected jobs. The system completes the outstanding reads and writes (up to 128) before suspending. This may take several minutes.
Restart	Restarts suspended jobs
Abort	Terminates the currently selected jobs. It immediately cancels all outstanding reads and writes. You can't restart the job, but you can resubmit it.
Show Log	Displays as much information as is written to the log so far. Use Exit on the File menu to close the log and return to the main window.
Select All	Selects all jobs in the Current Jobs display
De-Select All	Deselects all selected jobs in the Current Jobs display

Note

You cannot suspend some background jobs if you have MSAR libraries in backup mode. These include background jobs that update surfaces including background jobs associated with optical libraries.

Importing Documents


When you import documents, you move the indexes and other related information from incorporated media to the index database of the destination system. You can then retrieve and display the imported documents using normal access methods.

Note

In FileNet P8 Content Federation Services, the document class was created without indexes, there will be no index information on the media to import.


Before beginning the import operation, use Storage Library Control's Media Space Usage report to determine the surface ID of your original media and the number of documents written to that media (see "Media Space Use" on page 466).

This information is necessary to verify that the correct number of documents are successfully imported for each surface. Write down the information or print a copy of the report to use when verifying. See "Verify Document Import" on page 386.



Copy Annotations from Database to Media

To copy annotations from database to media, select Copy Annotations from Database to Media from the New menu. The following dialog box opens:



When you've made all your entries and selections, click OK to begin copying.

The following table shows the options for copying annotations from database to media.

Options for Copying Annotations from Database to Media

Option	Description
Surface ID of documents	Enter surface ID for documents whose annotations are to be copied.
Copy annotations to the surface containing the documents	Select this radio button to copy annotations to the surface containing the associated documents.

Options for Copying Annotations from Database to Media

Option	Description
Copy annotations to a different surface	Select this radio button to copy annotations to a surface other than surface containing their associated documents.
Destination Family Name	If you selected Copy annotations to a dif- ferent surface , use the pull-down list to select a family name.
Copy annotations from both sides of the media	Select this radio button to copy annotations for all the documents on this media.
Copy annotations from one side of the media	Select this radio button to copy only the annotations for the documents on the specified surface.

Importing Documents

Incorporate Media

Before you can import documents on media that came from another system, you must incorporate the media into your system. This process gives your system necessary information about the media (system serial number and surface ID from the original system).

If your system is not compatible with the foreign system, the media acquires a new surface ID when you incorporate it. A system is incompatible if the document ID and the surface ID ranges on the two systems overlap in any way.

Once the system can address the foreign media, you import documents from the media. Importing consists of copying information about the documents (index values, document classes, location of both the primary and transaction log copies, etc.) into the index and permanent databases.

Media incorporation actually happens in the foreground, so you cannot perform other administrative operations while you are incorporating media.

Media incorporation will not be allowed when a sub-set of the MSAR libraries on a server are in backup mode.

Background Job Control provides two methods to incorporate foreign media: one for Storage libraries (including rapid changers and MSARs) and one for ODUs.

During an incorporation, the local surface ID displays, but no information displays in the main window, nor does the job show up in the completed jobs display.

Storage Libraries

To load foreign media into a storage library (including rapid changers and MSARs):

1 Choose Incorporate Foreign Media from the Incorporate menu. The following dialog box appears.



- 2 Insert the Media. (Not applicable to MSARs).
- Select the media family name. The media family name does not need 3 to match the one previously used. However, for primary media, you must use a primary family name. For transaction log media, use a transaction log family name.

For an optical storage library

The family list includes all optical families, regardless of primary and tranlog types. Select the library and the media family name from the pull-down lists.

For an MSAR library

Only families with the same media type and primary/tranlog family type will be listed. Select the library from the pull-down list. Enter the filename in the text field or click the Browse button. If you click Browse, the filename you provide in the Select MSAR Surface Data File to Incorporate into Library dialog will appear on the Incorporate Foreign Media Setup dialog box.

Note

When the Select MSAR Surface Data File to Incorporate into Library dialog is first invoked, the dialog will use the standard "out_of_box" directory for the Image Services system as its starting point for the browse. If it is invoked again to incorporate another surface file, it will begin browsing in the directory the user last opened. This "last used" directory setting is lost when the Background Job Control application restarts.

4 Click OK.

- If you selected a media family that is incompatible with the media type used, a message box opens, stating the media type is incompatible.
- If you are using a rapid changer, a rapid changer dialog box appears. Enter the number of the slot where you inserted the foreign media and click OK.

The system inserts the media in a drive.

The system reads the original system serial number and surface number written by the exporting system, and displays them in a separate window.

- 5 Write down the original system serial number and surface ID (you need this information when importing the documents).
- 6 If this is the correct media, click OK to start the incorporation.
- The program displays the local surface ID, along with details. 7

(The following does not apply to MSARs) This process incorporates both sides for local and foreign media. After incorporating the first side, the system displays the following:

Incorporating the other side...


Note

(The following does not apply to MSARs) When you insert a 2X media into an 8X drive, the system enables the 2X media for reads only. If you later want to write to that read-only 2X media in a 2X or 4X drive, you have to manually enable the 2X media for reads and writes.

Optical Disk Unit

To manually load foreign media into an ODU:

1 Choose Manually Incorporate Foreign Media from the Incorporate menu to display the following dialog box.



- 2 Insert the media.
- 3 Select the library (ODU), drive number, and media family name from the pull-down lists.

The media family name does not need to match the one previously used. However, for primary media, you must use a primary family name. For transaction log media, use a transaction log family name.

4 Click OK.

> If you selected a media family that is incompatible with the media type used, a message box opens, stating the media type is incompatible.

> The system inserts the media in a drive and reads the original system serial number and surface number written by the exporting system, displaying these in a separate window.

- 5 Write down the original system serial number and surface ID (you need this information when importing the documents).
- 6 If this is the correct media, click OK to start the incorporation.

The program displays the local surface ID, along with details. After incorporating the first side, the system displays a pop-up window with the prompt:

Please flip the media and click OK.

7 Remove the optical disk, flip and reinsert it in the drive, then click OK.


Note

When you insert a 2X media into an 8X drive, the system enables the 2X media for reads only. If you later want to write to that read-only 2X media in a 2X or 4X drive, you have to manually enable the 2X media for reads and writes.

Import Documents from Media

See <u>"Importing Documents" on page 374</u> for a list of the steps and information required for importing documents from media. If you run more than one import job, you may want to note which job number goes with which surface ID. The local surface ID is the newly assigned surface ID on the importing system.

To import documents from media, select Import Documents from Media from the New menu to display this dialog box.



When you've made all your entries and selections, click OK. The following table describes the Import Documents from Media options.

Option	Description
Original System Serial Number (SSN)	Enter the original system serial number from the Incorporate Media option (see "Incorporate Media" on page 378).
Original Surface ID	Enter the original surface ID that the Incorporate Media option displayed. If you only want to import documents from one surface, enter that surface ID. For example, if the program displayed surface 4000 and you only want the documents from side B, you can enter 4001 instead of 4000.
Import All Documents	Select this radio button to import all documents from the specified surface or the entire media.
Use Document Headers File	Select this radio button to specify documents to import. If you use this selection, you can only import documents from one surface.
Document Headers File Name	Use the pull-down list to select a document headers file and import only the documents listed in that file. To create a document headers file, see "Create Document Header File" on page 453.
Import documents from both sides of the media	Select this radio button to import documents from both the specified surface ID and the opposite surface.
Import documents from one side of the media	Select this radio button to import documents only from the surface you entered as the original surface ID. You must select this option to import documents listed in a document headers file.
Are consistent document classes required?	Choose Yes to require that document class indexes be identical on the exporting and importing systems. Choose No to import the documents even if the classes have different indexes.
	If you choose Yes when differences exist in the document classes, the event log reports a message such as "The specified index does not exist." The system does not import the document unless it finds a corresponding record in the index database prior to the import. When it detects a document that existed prior to the import in the locator database but not in the index database, the system deletes it from the locator database if an error occurs during the import.
High priority?	Choose Yes to run the import job at the same priority as high priority reads (retrievals). Choose No to run the job at the lowest priority.

To verify that your documents were successfully imported, you must know the surface ID of the imported media and the number of documents written to that media. Use Storage Library Control to get this information from the original media (see "Media Space Use" on page 466).

To verify your import operation:

- 1 Run Storage Library Control and choose Media Space Usage from the Library Configuration window's Reports menu (see "Media Space" Use" on page 466 for details).
- 2 Compare the number of active documents on the destination media with the number of active documents on the source media.
- 3 If the number of active documents on the destination media does not agree with the number of active documents on the source media, use any text editor to open the file located in the file path for your operating system:

/fnsw/local/logs/bkglog/lmpLog.xxxxxx (UNIX)

where xxxxxx is the job number.

The following example shows the contents of an import log named ImpLog.000052.

```
Information for import job number 52 started at Thu Oct
29 17:00:57 2003
Job parameters:
Input surface id: 3056
Import from both sides: no
Import all files: no
Document class must match exactly: no
Reset surface statistics: no
Insert into doctaba: yes
Security options: from document class
Update scalar numbers table: no
Update WorkFlo queue: yes
Skip import of deleted documents: no
High priority: no
Input media filename: FN DESCRIPTOR 19980925 172206
End of job information:
Number of errors encountered: 0
Number of documents read from input media: 1441
Number of documents imported: 1441
Number of duplicate documents in docs table: 1441
Job completed at Thu Oct 29 17:02:10 2003
```

Look at the log for the surface IDs of the media from which you 4 imported documents.

This log lists any errors that occurred during the import operation and the document IDs that were not imported.

5 If the system failed to import any of the documents, you can reimport them.

You don't need to use the Incorporate Foreign Media option again, as the importing system already recognizes the media. If the documents were not imported due to some easily-remedied error (for example, the document classes were inconsistent), then fix the problem and reimport the documents.

The stdocimp command line utility gives you a little more flexibility in document importing: you can skip over documents, use a file list, or both. See "stdocimp" on page 530 for details on using this utility.

Copying Documents


You can copy all documents from a specified media surface to a specified media family.

You can copy particular documents to a specified media family by designating a document headers file.

You can also copy particular documents from storage media to a media family based on a list of document ID numbers in a file. See "Copy Specific Documents from a Surface to a Family" on page 393.

Copy All Documents from a Surface to a Family

To copy all documents from a surface to a family, select Copy Documents from the New menu to display this dialog box.



When you've made all your entries and selections in the Copy Documents dialog box (see table below), click OK to begin copying.

Copy Documents dialog box

Option	Description
Surface ID	This field requires a media surface ID number.
	To copy from only one surface, enter the correct ID for that surface (on two-sided media, you can copy a single surface only if you're using a document headers file).
	To copy from both surfaces, enter either surface ID number.
Use Document Headers File	Select this radio button to limit the documents copied. Use the pull-down list to select a document headers file name. To create a document headers file, see "Create Document Header File" on page 453.
Copy All Documents	Select this radio button to copy all non-deleted documents on the specified surfaces.
Destination Family Name	Use the pull-down list to select the name of the media family to which you are copying the documents (the family does not have to match the source media family).
	The system copies documents to the current write surface of the family you specify. If that surface does not have enough space, the system allocates a new surface ID and requests new media through an RSVP message (see "RSVP Messages" on page 428).
Copy documents from both sides of the media	Check this box to copy both sides of the media. If you entered a document headers filename, you can only copy from one side of the media. This option is ignored for MSAR surfaces.
Copy documents from source media	Select this radio button to copy from input surface ID.
Copy documents from alternate media	Select this radio button to copy from alternate media If the input surface is a primary media, the alternate is the transaction log. If the input surface is a transaction log, the alternate is the primary media.
Both primary and alternate copies of the documents are available	Check this box if both the primary media and the transaction log media are in the storage library. If a read error occurs on the first copy accessed, the system will try the other copy when you indicate that both copies are available.

Copy Documents dialog box, Continued

Option	Description
Update document indexes with current information	Check this box to replace document indexes on the storage media with values from the index database on magnetic media.
	Since index values can change on magnetic media (where DOCTABA resides), you can either copy the more current information from the index database or copy the original values that were written to storage media when the documents were committed.
Update database to refer to new document locations	Check this box to update the permanent database to look at the destination media for the copied documents (the documents on the original source surface become inaccessible).
Copy annotations	Check this box to copy all annotations associated with the documents being copied.
	Copying annotations is not necessary unless you are moving the copied documents to another system where the annotations must also be available.
Generate log of	Check this box to create a log showing which documents were copied.
documents copied	The copy log file for UNIX platforms is: /fnsw/local/logs/bkglog/CPYDoc.nnnnnn
	The copy log file for Windows Server platforms is:
	<pre><drive>:\FNSW_LOC\logs\bkglog\CPYDoc.nnnnnn</drive></pre>
	(nnnnnn is the document copy job number). See "Completed Jobs" on page 416.

Note

The destination family list will display families that are write compatible on the local server only. All family types (primary and tranlog) are displayed so that users can copy from a primary surface or vice-versa. A pop-up window appears if the family type does not match the destination family and the user can then choose whether to continue the write.


Copy Specific Documents from a Surface to a Family

You can specify certain documents to copy from storage media to a media family by creating a file. The file you create must consist of a list of document IDs with one ID per line. Each line is a string of ASCII digits followed by a newline character. Spaces and tabs are ignored.

Create a Document File List

You can use one of two methods to create a document file list.

- Generate a list of document IDs by creating a file of document IDs and placing it in the fnsw temporary directory.
- Use Find Open Documents to generate a list of open documents.
 - Select Find Open Documents from the New menu to display this dialog box.




Enter the surface ID of the media you are interested in.

c Enter a filename for the list of open document numbers. The file automatically goes to the fnsw temporary directory.

For example:

/fnsw/local/wfl/tmp (UNIX)
<drive>:\FNSW_LOC\wfl\tmp (Windows Server)


If you do not enter a name, the system does not create a file. Instead, it generates a report showing how many documents of each status (open, closed, deleted) are on the media. To see the report, choose Results of Find Open Documents from the Completed menu (see "Completed Jobs" on page 416 for details).



d Click OK to start the program or click Cancel to close the window without starting the program.

Copy Documents Using a File List

To copy documents using a file list, select Copy Documents Using File from the New menu to display this dialog box.



When you've made all your entries and selections (see table on following page), click OK to begin copying.

The following table describes the options in the Copy Documents Using File List dialog box.

Copy Documents Using File List options

Option	Description
File containing document numbers	The file must be in a specific directory which is already entered in the text field. Click the Select button to display a standard file selection dialog box and choose a filename.
Destination Family Name	Use the pull-down list to select a media family name to copy the documents to (the family does not have to match the source media family). If the current write surface of the family you specify does not have enough space, the system allocates a new surface ID and requests new media through an RSVP message (see "RSVP Messages" on page 428).
Copy documents from primary media	Select this radio button to choose the primary media as the source media.
Copy documents from alternate (tranlog) media	Select this radio button to choose the transaction log media as the source media.
Both primary and alternate copies of the document are available	Check this box if both the primary media and the transaction log media are in the storage library.
Update document indexes with current information	Check this box to replace document indexes with values from the index database.
Update database to refer to new document locations	Check this box so the permanent database looks at the destination media for the copied documents (the documents on the original source media will be inaccessible).

Copy Documents Using File List options, Continued

Option	Description
Copy annotations	Check this box to copy all annotations associated with the documents on the media. Copying annotations is not necessary unless you export the copied documents to a system where the annotations must also be available.
Generate log of documents copied	Check this box to create a log showing which documents were copied (see "Completed Jobs" on page 416). The copy log file is:
	The copy log file for UNIX platforms is:
	/fnsw/local/logs/bkglog/CPYDoc.nnnnnn
	The copy log file for Windows Server platforms is:
	<pre><drive>:\FNSW_LOC\logs\bkglog\CPYDoc.nnnnnn</drive></pre>
	(nnnnnn is the document copy job number). See <u>"Completed</u> <u>Jobs" on page 416</u> .

Consolidating Media

Consolidating media is a variation of the copy operation so the same considerations apply regarding priority, magnetic disk space needed, automatic restart, pending writes, etc.

CAUTION

The Consolidated Media function does not copy updated index values from the index database. Instead, it copies original index values from the media. Because of this, it is best to use this function to copy documents if index values may have been changed. It is recommended that you use **Copy Documents** with the **Update document indexes with current information** option instead. If the document indexes have been updated, the updated document indexes get copied over. Note that this has slower performance but the correct indexes are copied with the document.

When the number of active documents is too low to justify keeping particular media in the library, that media is a candidate for consolidation. To decide when to consolidate media, refer to the Media Space Usage report (see "Media Space Use" on page 466).

You can use two methods to consolidate media.

 The Consolidate Media function in Background Job Control's New menu assumes you do not want to retrieve the outdated documents again and copies all documents that are not deleted to other media. Then you can remove the old media to make more room in the storage library.

If any pending writes exist for the media being consolidated, the system automatically switches those pending writes to other media when you start the consolidation. This method copies closed documents, so you must delete any closed documents you don't want copied before consolidating.

An alternative is to use the Find Open Documents option to generate a list of open documents (open documents have been neither deleted nor closed). Then you copy the open documents to other media and remove the source media (without using the Consolidate Media operation). If you consolidate in this way, you do not need to delete documents to consolidate media. The Background Job Control interface does not accept surface IDs greater than 999999. For surface IDs larger than 999999, consolidate media using the stdoccpy command.

Consolidate Media


Before using Consolidate Media, you can first use the **deldocs** tool to delete all the closed documents you do not want to copy. Deleting documents means you are deleting records from the index and permanent database (and cache, if you have no storage media). Use this command to delete all the documents on both surfaces:

deldocs -s <surfaceID> -b

Enter the surface ID. The **-b** option specifies both sides. See "deldocs" on page 492 for more information about the deldocs tool. The deldocs tool does not erase the media. You may use Consolidate Media to erase the media (see Step 3 below). If no documents exist for the consolidate operation to copy, the media is simply erased. See "Erasing Media" on page 406 for another way to erase media.

To consolidate media:

1 Select Consolidate Media from the New menu to display this dialog box.



Enter the surface ID of the media you want to copy. The system con-2 firms that the ID is valid. If you are going to copy documents from only one surface, be sure to enter the correct surface here. The system consolidates both sides of the media even though you specify only one surface ID.

Be aware that not all the families are listed in the Destination Family Name pull-down list. Only families that are write-compatible and with a matching primary/tranlog type are listed. Select the appropriate media family name from the Destination Family Name pull-down list.

The destination family is determined in one of the following ways:

- If the source surface has clustering, only the original family will be displayed.
- If the source surface does not have clustering, the following rules will be used to determine the list of eligible families, and the program writes to the current write surface of the family selected:

- For non multi-OSAR servers, it will display all write-compatible families with matching primary/tranlog family types will be displayed.
- For optical surfaces on multi-OSAR servers, only the original family will be displayed.
- For MSAR surfaces (including MSAR converted surfaces) on multi-OSAR servers, the families which match the following criteria will be displayed:
 - •Families have matching primary/tranlog family type.
 - Families are write compatible.
 - •Families are only on the local server.
 - All write surfaces have a preferred library.
- The system copies the documents to the current write surface for this family.
- 3 Check the Erase Media box to erase the source media.

CAUTION

Do not stop the Image Services software while an erase operation is in progress. If you shut down the Image Services software in the middle of an erase, the erase procedure resumes when you recycle the software, making the drive unavailable for any other purpose (regardless of priority) until the erase is complete.

4 The Update document indexes with current information option is checked by default which means the document indexes WILL be updated. If the option is unchecked, the document indexes will NOT be updated. The original index values entered on the storage media may have been changed in the index database. While copying documents,


you can have the system insert the most recent index data instead of copying the original information from the media. To do this, the "Update document indexes with current information option must be selected.

- 5 Click OK to begin consolidation.
- 6 This step does not apply to MSAR. If you're erasing media, after consolidation eject the media by surface number, then reinsert the media (see "Eject Media by Surface ID" on page 446 and "Insert Media" on page 441 for details).

Find Open Documents

To consolidate media (open documents only) without deleting closed documents, use the following procedure:

Select Find Open Documents from the New menu to display: 1



- 2 Enter the surface ID of the media you are interested in.
- 3 To create a file listing the IDs of open documents, enter a filename.
 - The file automatically goes to the /fnsw/local/logs/bkglog (UNIX) or <drive>:\FNSW LOC\logs\bkglog (Windows Server) directory.
- 4 Click OK.

- 5 Select Display Results of Find Open Documents from the Completed menu to see the number of open, closed, and deleted documents on the media.
 - The report also shows the total number of documents and what percent are open (see "Results of Find Open Documents" on page 418).
- Submit the list of open documents to the Copy Documents Using File 6 List option, using the filename you entered in Step 1 above (see "Copy Documents Using a File List" on page 395 for details).

Rebuilding Media

The system restores information from lost or damaged media by reading the documents from the alternate copy of the media. For primary media, the alternate is the principal transaction logging media; for transaction logging media, the alternate is the primary media. The rebuild operation copies all documents on both surfaces of the media you are rebuilding, even though you enter one surface ID.

After reading the documents, the system writes them to the current write surface as described in "Consolidate Media" on page 399. If the source media has pending write requests, the system automatically switches them to the current write surface when you start the rebuild.

Select Rebuild Media from the New menu to open the Rebuild Media dialog box.



Enter the surface ID of the media to rebuild and click OK.

Note

The rebuild operation copies to the current write surface, so the result is not a duplicate copy of lost or damaged media. Rather, documents for which one copy was previously unavailable are now available from both their primary and transaction log families.

CAUTION

For converted MSAR surfaces, rebuilding this type of surface will result in writing it to its originally assigned, optical disk family. If you don't want this to happen, you need to use the **stdoccpy** utility to rebuild the

converted MSAR surface. The **stdoccpy** utility will allow you to specify an MSAR document target family.

The following syntax examples are for surfaces that **DO NOT** have clustering turned on:

To rebuild a primary surface from a tranlog surface:

stdoccpy -family <familyname> -surface <surfaceid> -bothsides -updatedb -altsurf -onecopy -findby db

To rebuild a tranlog surface from a primary surface:

stdoccpy -family <familyname> -surface <surfaceid> -bothsides -updatedb -onecopy -findby db


Erasing Media

If you use erasable media as a temporary backup, you may erase the media when the information is no longer needed, even if the media contains open documents.

CAUTION

This procedure entirely deletes the data from the media. Be sure the media you are erasing contains information you no longer need.

1 Select Erase Media from the New menu.



- 2 In the dialog box, enter the surface ID of the media surface you want to erase.
- 3 Click OK.

CAUTION

You cannot interrupt the media erase process. Interruptions during erase disk may destroy disks or cause the system to lose track of disks. Any drives performing an erase are unavailable for any other purpose (retrievals, committals, etc.) until the erase is complete.

Note

Erasing media on systems with a large permanent database can be a lengthy process because a full database search is automatically performed to find all documents on the media. The larger the database, the longer this search may take.

Phases for Erase Media

Erase Media has two phases:

Phase 1: Generates a document ID list from the source surface. If the corresponding document is found in the DOCS Permanent database, a delete or update of the database is performed. When this phase is done, the background job immediately forwards to Phase 5.

Phase 5: Erases the source media.

For Optical Disk:

The surface is erased, becomes unlabeled and then is unloaded to a slot. The surf_info entry and surface locator entry (if it's a multiple storage library server configuration) are deleted.

For MSAR Surface:

The surface is marked "out-of-box" and is removed. The surf info entry and surface locator entry (if it's a multiple storage library server configuration) are deleted. The lib_surfaces entry is also deleted.

Note:

There are no Phases 2, 3, or 4 for Erase Media.

Erasing Media During Consolidation

To erase media during Consolidation

- 1 Check the Erase Media check box in the Consolidate Media dialog.
- After the consolidation, eject the media by surface number. 2
- 3 Reinsert the erased media.

Erasing Media Without Consolidating

If you use erasable media as temporary backups, you may erase those media when the information on them is no longer needed, even if the media contain documents (open or closed).

To erase media without Consolidating

- 1 Select Erase Media from the New menu.
- 2 Enter the ID of the surface you want to erase and click OK.

CAUTION

This procedure entirely deletes the data from the surface. Be sure the surface you are erasing contains information you no longer need. Also, be aware that the entries associated with documents on this media are being deleted from the Index database (doctaba table) and the MKF databases (docs table)

Migrating Documents

In a cache-only system, committed documents are stored on magnetic disk. When you add a storage library or ODU to a cache-only system, you can use the Migrate Docs option from Background Job Control's New menu to move those documents to your new storage media.


Note

Migrate Docs does not support documents dated earlier than 1970.

When converting a cache-only system, you must verify that **all families** are changed to supported media. Once a single write has gone to an optical or MSAR surface, any committals to an unsupported media will end with an invalid media error. It is not possible to support families that write to media and others that will continue to write to the cache-only, as they did before creating a library. For example, cache-only system has families that write to the non-existent optical media, and once the MSAR library is created, all families must be altered to write to MSAR media or all doc classes must be altered to write to families which in turn write to supported media.

To migrate documents to storage media:

1 Choose Migrate Docs from the New menu to display the Migrate Document dialog box.



- Select a migration option by checking the appropriate radio button: 2
 - Check All unmigrated documents to migrate all committed documents on magnetic disk to storage media.
 - Check All unmigrated documents older than create date (<format>) to migrate all committed documents older than a specified creation date. Enter the desired creation date in the text field using the format specified for your operating system.

The system uses the date format configured through the Image Services server's operating system. For example, while a server installed in Europe might use the dd/mm/yyyy format, a server installed in the United States most likely uses the mm/dd/yyyy format.

- Check All documents using file list to migrate all committed documents using a specified file list. Click the Select button and, in the dialog that appears, choose the name of the file list for the documents to migrate.
- 3 When you have made your selections, click OK to accept the migration choices or Cancel to exit the dialog box without making any changes.

Converting Optical Surface to MSAR Surface

Converting optical surface to MSAR aids users in deploying MSAR libraries by providing the capability of doing a fast copy of an optical surface to an MSAR surface. The only data needed to process the conversion are the optical surface ID and the MSAR library the surface is going to.

Note

In a cache-only system, you can move your committed documents to an MSAR library using Background Job Control's procedure "Migrating Documents" on page 409.

Note

All foreign surfaces that are converted to MSAR will gain the local system serial number, in effect becoming local surfaces.

When the convert background job starts, the source optical is disabled for writes. Note that after the docs table has been updated, if the **Update database to refer to new locations** option is selected, the source optical surface will also be disabled for reads. If this option is not selected, the source optical surface will not be disabled for reads.

If the Convert both sides of media option has been selected, and the source optical surface is a two-sided media with documents on both sides, two MSAR surfaces will be created from a single conversion job. The new MSAR surfaces will be single-sided media and the surface ID will always be an even number.

If the source optical surface is blank, or if it's labeled but has no documents, the conversion will skip the side.

Conversion Phases

There are five distinct phases when converting an optical surface to an MSAR surface:

Phase 1 The lib surfaces entry is created. The MSAR surface file location is specified.

Phase 2 The optical surface is read and the MSAR surface is written.

Phase 3 The MSAR surface is inserted into an MSAR library. The surf info, surf locator, and lib surfaces tables are updated. OSA and SRF shared memory are also updated.

Phase 4 The docs table is updated, if specified.

Phase 5 The job completes.

When the convert background job enters Phase 4 (update docs phase), the job cannot be aborted using the **Abort** button.

To convert optical surfaces to MSAR surfaces:

- 1 Choose Convert Optical Surface to MSAR from the New menu to display the Convert Optical Surface to MSAR Surface dialog box.
- 2 Enter the Surface ID of the optical surface you want to convert in the Surface ID field.
- 3 Select the MSAR library in which you are placing the newly converted surface from the MSAR Library pull-down.
- 4 Select the OK button.

Note

Any surface converted from optical to MSAR becomes read-only. You can no longer write to that particular surface. You need to copy annotations to the source optical surface or update the document headers BEFORE running the conversion. Also, to prevent any new write requests from being committed to the optical family after an optical to MSAR conversion, make sure the document classes that previously belonged to the optical family are redirected to the appropriate MSAR family BEFORE beginning any conversion tasks.

The Optical to MSAR conversion process generates an MSAR surface file and a lib_surfaces entry. The following table lists a number of actions that would cause the conversion to terminate prematurely and indicates whether the utility would delete these two key files.

Conversion Termination Cause	Comment
Error occurred during the conversion process	In addition, any entries in the surf_info or surf_locator tables will also be removed.
Recycling of Image Services Software or msar_convert_bkg is terminated during the conversion	The MSAR surface file and lib_surfaces will not be deleted and the conversion will not automatically restart.
Target MSAR directory runs out of space	The MSAR surface file and lib_surfaces will be deleted and the conversion will be terminated.
Errors reading optical disks	The MSAR surface file and lib_surfaces will be deleted and the conversion will be terminated.
	Although lower layers of OSAR Services (ODX and ODU) do retry's, media errors (ODX_zbaddata error tuple) that may be caused by a dirty optical disk cause the conversion to terminate. The user will also be instructed to eject the particular optical disk, clean the disk and retry the conversion command.

Whenever possible, the conversion program skips unwritten sectors of the optical disk.


While an optical surface is being converted to an MSAR surface, the associated MSAR surface will be marked as out-of-box. This is necessary so that the surf_file field will be created in the lib_surfaces table.

Generating Reports

Background Job Control uses the log files to create the lists displayed in the Completed Jobs and Results of Find Open Documents reports.

Completed Jobs


Choose Completed Jobs from the Completed menu to display the Completed Jobs report.



Completed jobs are listed by job number. The type of job (import, copy, etc.) appears in the second column, with the completion time in the third.

To delete background job logs, select the logs to delete and click on Delete Log.


Select a job and click the Show Log button to display a summary of the job parameters and errors or other messages (see also "fn_msg" on page 500).



Results of Find Open Documents

This report shows the results of the Find Open Documents function (see "Find Open Documents" on page 402). Open documents have been neither closed nor deleted. This report shows the number of open, closed, and deleted documents on specified storage media and the percent of open documents.

After running the Find Open Documents function, select Results of Find Open Documents from the Completed menu to display this report.



Use the horizontal scroll bar to see the last two columns.

When your library is nearly full, you can either consolidate media (which copies open and closed documents to another surface) or you can copy just the open documents. See "Consolidate Media" on page 399 and "Copy Documents Using a File List" on page 395. Then you can remove the old media to make another slot available in the library.

Storage Library Control

To communicate with a storage library, you must run Storage Library Control (SLC) at the server that controls the library. This server is called the storage library server.

If you have more than one storage library server, you need to run SLC at each server where libraries are being used. On operating systems that support X-windows, you can remotely log in and run SLC using the X-window protocol.

SLC displays informational messages sent by a storage library. An informational message tells you about an event that has already occurred. You do not need to respond to these messages.

An RSVP message can come from any device and requires the operator to respond and perform some action, such as loading media.

In addition to responding to messages sent by SLC, you can initiate operations such as enabling or disabling drives, slots, etc.

You can use SLC to:

- Add media to and remove media from a storage library or an optical disk unit (ODU).
- Enable or disable one or both surfaces of storage media for reads, writes, or both.
- Preformat storage media to reduce the risk of loading errors on ODUs. (Does not apply to MSAR.)

Media

Media refers to any material on which data is stored (magnetic disk, optical disk, magnetic tape). SLC is concerned primarily with permanent storage media, usually optical disks or MSAR systems.

Preformatted

If you use an optical disk unit (ODU), which requires manually inserting storage media directly into the drive, you should preformat your storage media as soon as you receive them (see "Preformat Media" on page 440). Preformatting marks the surfaces of storage media as Side A and Side B. SLC assigns surface IDs when the media are first written, not during preformatting.

Preformatting should be done by an operator trained to correctly insert unformatted storage media. If you do not preformat your storage media, be sure all operators know how to load blank storage media properly. Failure to load blank storage media properly into an ODU can make the media partially unusable.

Surfaces

All storage media [except Plasmon (Philips) 12" optical disks] and MSARs have two surfaces, Side A and Side B. The system assigns an even numbered surface identifier (for example, 3000) to Side A and the next odd number (for example, 3001) to Side B. Surfaces are automatically selected when you specify a media family. In cases where you explicitly name a surface ID, the system ignores the disabled status of a surface.

For example, you can copy documents from a surface that has been disabled for reads because you are specifying the surface ID.

When copying documents to storage media, you specify a family name, not a surface ID. In this case, the software does not select a surface that has been disabled for writes; it selects other storage media in that family.

Transaction Log

You may need to change the transaction log media type if you import documents from a system where the transaction log media type is different from that used on your system (see "Change Media Type" on page 450). While the system is importing documents from the primary media, it determines whether an alternate copy of the document exists on transaction log media. If so, it can obtain the media surface number corresponding to that document.

On systems that use different media types for primary families, we recommend that you create matching transaction log families for each media type (see "Media Families" on page 86). However, if you only want one kind of transaction log, regardless of the types of primary media, then you may need to use the Change Media Type function.

To import documents from transaction log media later, you need to know the media's surface IDs and you must correct their sizes before inserting the media into a storage library. Since all media in a family (including transaction log families) must be the same size, the Change Media Type function changes the foreign transaction log's media type and marks it as not writable.

Document Header Files

At least one document header file is assigned to all storage media. This file contains a record of the documents committed over a particular period of time. The system creates the first file; you create additional files as needed (see "Create Document Header File" on page 453). When you create a new document header file, the file contains a record of the documents committed from that point until you create the next document header file.

If you are entering and committing documents on one system for export to another system, use document header files to group documents chronologically. You can then take the storage media to another system and import only the documents listed in the specified document header file.

Storage Media Insertion (non-MSAR only)

The following tables specify which direction Side A of optical disks should face when inserting them in the storage library or ODU.

Storage Media Insertion Guide

FileNet OSAR, Hitachi 12" Disk Drive (2.6 GB, 7 GB)

New Media/Even Surface	Side A must face to the right
Odd Surface	Side A must face to the left

Hitachi ODU, 12" Disk Drive

New Media/Even Surface	Side A should face down
Odd Surface	Side A should face up
The unit reads or writes only the side facing down.	

HP or IBM 5" Library (Autochanger), 5" Disk Drive

New Media/Odd Surface	Side A should face up
Even Surface	Side A should face down

Hitachi 5" Library (MOSAR), 5" Disk Drive

New Media/Odd Surface	Side A should face up
Even Surface	Side A should face down
Document Services requires Side A face down on new media and writes that side first.	

HP or IBM ODU, 5" Disk Drive

New Media/Even Surface	Side A should face down
Odd Surface	Side A should face up
Document Services requires Side A face down on new media and writes that side first.	

Storage Media Insertion Guide, Continued

New Media/Even Surface	Side A should face up
------------------------	-----------------------

FileNet OSAR-40/50HTL/S, Plasmon (Philips) 8000 Disk Drive

New Media/Even Surface	Side A must face up
------------------------	---------------------

FileNet OSAR-107/123/144GTL/S, Plasmon (Philips) 6000 Disk Drive

New Media/Even Surface	Side A must face to the left
------------------------	------------------------------

FileNet OSAR-107/123/144HTL/S, Plasmon (Philips) 8000 Disk Drive

New Media/Even Surface	Side A must face to the left
------------------------	------------------------------

Plasmon (Philips) RapidChanger, Plasmon (Philips) 6000 Disk Drive

New Media/Even Surface	Side A must face to the right
------------------------	-------------------------------

Plasmon (Philips) RapidChanger, Plasmon (Philips) 8000 Disk Drive

New Media/Even Surface Sid	de A must face to the right
----------------------------	-----------------------------

Plasmon (Philips) ODU, Plasmon (Philips) 6000 Disk Drive

New Media/Even Surface	Side A should face up
------------------------	-----------------------


Plasmon (Philips) ODU, Plasmon (Philips) 8000 Disk Drive

New Media/Even Surface	Side A must face up

Starting Storage Library Control

After you are logged on to the Image Services software, select Storage Library Control from the Application Executive's Applications menu.

The SLC main window displays. The title of the window shows the name of the local storage library server.



The SLC main window has four sections:

- Informational Messages section displays status information.
- RSVP Messages section displays operator instructions.
- The RSVP/INFO Trigger field indicates whether or not the RSVP/ INFO triggering script is enabled.
- Storage Libraries section displays a line for each library or ODU attached to the server.

Messages

SLC displays messages in chronological order, oldest to newest. By default, the list scrolls to the newest item. You can change this by turning off New Message Scrolling in the Miscellaneous menu. The system updates messages at specified intervals, but you can update them immediately by choosing Refresh from the Miscellaneous menu.

To see the rest of a message that does not fit within the window, double-click the message (or select a message and click the Info or Show button). In the extended message window that appears, click the Next and Previous buttons to display information about other messages.

Informational Messages

Informational messages report system events that have already occurred. These messages don't require any action on your part. Messages accumulate until you delete them, unless you set a limit (see the table under "Message Display" on page 430).

- To delete informational messages, select them and click the Delete button.
- To select more than one message, drag through a group of consecutive messages, or select the first one, then hold down the Shift key while selecting the last one.
- To select nonconsecutive messages, hold down the Control key while selecting.
- To deselect any previously selected message, click on any message without pressing one of these keys

RSVP Messages

An RSVP message requires some action or response on your part. For example, an RSVP message might instruct you to load media. Doubleclick on a message (or select a message and click the Info button) to see a dialog displaying the entire message and the time it arrived. To respond to an RSVP message, select the message and click Reply. This displays a dialog box that contains any additional messages associated with the RSVP and tells you what to do.

For MSAR insertion RSVP, if there are multiple MSAR libraries configured, you will first see a dialog box which asks you to select the MSAR library. Then you see the file open dialog box to select the MSAR surface file. (The requested MSAR surface ID is displayed in the dialog box heading.)

Always follow the messages exactly. For example, on an ODU, wait until SLC prompts you to remove storage media before you press the eject button. For ODUs, you always disable a drive before removing the storage media, unless you are following explicit prompts from SLC to remove media.

See the Online Help for a description of the RSVP messages. See MSAR Procedures and Guidelines for Image Services for more information about MSAR RSVP messages.

Normal-Mode RSVP Messages

In normal mode, the robotic arm in the storage library moves storage media between slots, drives, and the loading compartment (the I/O station). You see normal-mode RSVP messages when the robotic arm is operating normally. The normal-mode RSVP messages prompt you to insert or remove storage media.

Manual-Mode RSVP Messages

You see manual-mode RSVP messages on a system that has an ODU instead of a storage library. In manual mode, you must physically insert and remove storage media.

If you decide to eject media from an ODU without instructions from SLC, be sure to disable the drive (click Disable on the SLC main window) before pressing the eject button.

When inserting media in the drive without instructions from SLC, do not enable the drive until you have inserted the media. Clicking Enable causes SLC to read the label. Until it reads the label. SLC does not acknowledge the presence of the media.

You may also see manual-mode messages on a system with a storage library. Usually your service representative must respond to messages requesting manual operation of a storage library.

RSVP/INFO Trigger

This field indicates if a user-provided program/script to help with the response to an RSVP or INFO event is turned ON or OFF. If no script or program is configured, the field will simply say OFF. If a script or program is configured, the field will say ON and the configured script or program name (full path) will be displayed in the Script File field. The Script File field only appears when a RSVP/INFO Trigger script or program is configured.

A script or program is configured by filling in the appropriate information on the Scheduling sub-tab of the Server Application Services tab in fn_edit. This script or program will be launched automatically and, if called for in the script, an e-mail can be sent to the System Administrator when appropriate RSVP or INFO messages are received. For more information about using RSVP/INFO trigger, see "Appendix E -Message Triggering" on page 626.

Message Display

The Miscellaneous menu contains selections for changing the messages display.

Miscellaneous menu

Menu Selection	Effect on Messages			
Refresh	Select Refresh to update all information on the SLC main window immediately. This refreshes both informational and RSVP messages, as well as the storage library status.			
New Message Scrolling	Select New Message Scrolling to scroll list boxes to display newly-arriving messages (unless that would remove the currently-selected message from view). Deselect this option if you do not want the list box to scroll to the newest message.			
Options	Select Options from the Miscellaneous menu to set polling frequency, message beeping, and auto deleting. These options apply to the current session only; you must reset them each time you start SLC.			
	Polling Policy: SLC's default setting polls every 10 seconds for new messages from the libraries. To change the polling frequency, select the number in the box and replace it with another. By default, SLC checks less frequently for changes in status to libraries, drives, grippers, etc. Specify the number of polls to complete before checking status changes.			
	New Message Beeping: By default, the terminal beeps for each new message. If you have set beeping parameters at the operating system level (such as flashing the screen instead of beeping), SLC uses those conventions. To change the default, enter the number of RSVP and informational messages that should occur before the terminal sounds a beep. Enter 0 (zero) to turn off beeping.			
	Informational Message Auto Deletion: By default, messages accumulate until you delete them. To change the default, enter the number of informational messages to accumulate before the system automatically deletes the oldest messages. You can always manually select and delete messages.			

Message Automation

Argument	Field	Type	Valid Range	Comments
1	type	integer	1-2	1=RSVP
				2=INFO
2	error/event indicator in decimal format	Variable String [40]		

Storage Libraries

The Storage Libraries list at the bottom of the SLC main window contains one line for each storage library. This area describes each library type, how many slots it has, how many drives it has, whether it's enabled or disabled, and its operating mode (Normal is the operating mode for libraries and Manual is the operating mode for ODUs).

Enable/Disable Library

To enable or disable a library (this button is grayed out for MSARs), select the appropriate button in the Storage Libraries area. For more information see "Enable/Disable Library" on page 437.

Show Library Information

To display a summary of information for a given library, select the library and click Show. For more information, see "Show Library" on page 433

Backup Library

The Backup button in the Storage Libraries area is for MSARs only and allows the user to easily place MSAR libraries in an MSAR-read-only mode for backup purposes. This button is grayed out for optical libraries. This button brings up a simple Yes/No message box that asks "Are you sure you want to put MSAR Library X into Backup mode?" if it isn't already in Backup mode and "Are you sure you want to take MSAR Library out of Backup mode?" if it already is. If there is any problem updating the checksum on any in-box MSAR surface, transition to Backup mode will not succeed.

We recommend that you put all MSAR libraries into Backup mode at the same time. You cannot insert Media when a subset of the MSAR. libraries in a server are in Backup mode. You can, however, insert media when all MSAR libraries in a server are in backup mode or normal mode.

This restriction insures that MSAR surfaces are not moved from an MSAR library, in backup mode, to an MSAR library, in normal mode. The expectation is that ejecting and inserting MSAR surfaces is not done as often as ejecting and inserting optical surfaces. This restriction is maintained for insertion RSVPs associated with MSAR surfaces and the insertion Xslc menu.

Backup mode is persistent across an Image Services software restart.

Note

The request to enter or exit Backup mode is not processed by Storage Library Control, but rather it is gueued to be processed by the scheduler. Therefore, the request may not be processed immediately.

CAUTION


When a library is in Backup mode, **DO NOT** change any library's configuration. Doing so will invalidate the checkpoint file which is where the library mode is being stored. This will cause the library to lose the Backup mode after Image Services is recycled and go back to Normal mode. When the library is in Normal mode, all pending write requests will be processed and the surface files will be written to. This is undesirable when the surface files are still being backed up.

Library Management Functions

The storage library management functions described here are accessible from the SLC main window's buttons and menus or from the Library Configuration window's buttons and menus.

Show Library

In the SLC main window, select a library or ODU and click the Show button (or double-click on the library) to display the Library Configuration window for that library. Information in the Library Configuration window summarizes the state of the selected library, including information about the type of hardware, pending requests, and information about the library's drives, as shown in the following example.



Use the Library Configuration window to:

- "Enable/Disable Drive" on page 438
- "Calibrate Library" on page 437
- "Enable/Disable Slot" on page 438
- "Enable/Disable Grippers" on page 439
- "Insert Media" on page 441
- "Preformat Media" on page 440
- "Eject Media" on page 445

Library

This area describes the library type, how many slots it has, its operating mode (Normal is the operating mode for libraries and Manual is the operating mode for ODUs), and other self explanatory status fields. The MSAR Surface Directory field contains the directory location where all created MSAR surfaces will be placed for each MSAR library. If you need to, you can edit this field in the System Configuration Editor 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.

Pending Requests

The Pending Requests area of the Library Configuration window shows how many requests of each type are pending for the library, except surfaces in disabled drives or slots, surfaces in the process of being ejected, and surfaces waiting for identification.

The following table shows the types of library requests and what the values represent.

Type of Request	Value
Reads	Sum of pending high and medium reads (retrieval requests and read aheads for anticipated retrieval requests)
Writes	All writing to storage media except that done in the background (copy and import jobs started in Background Job Control)
Prefetches	Reads performed by programs (like WorkFlo) that gather data from storage media during off hours so the data is in cache for the next day

Type of Request	Value
Other Requests	Any read and write requests queued by copy and import jobs started in Background Job Control
Total Requests	Number of all pending requests

Drives

The Drives area of the Library Configuration window shows the number and type of each drive in the library.

For each drive, a line of information displays the surface ID of the media in the drive, the time the media was inserted in the drive, the number of read/write operations, number of errors that occurred on the current surface, whether the drive is configured to favor write requests, and the status of the drive (enabled or disabled).

This is also the area where you enable or disable drives. You might disable a drive if you see errors associated with the drive. In some cases, the system enables or disables drives for you. When you enable a drive, SLC reads and identifies any storage media in the drive, whether the storage library is in Normal, Manual, or Disabled mode. Your service representative might reserve a drive or an entire library for diagnostics. You cannot change a reserved state.

Message Display

The Library Configuration window's Miscellaneous menu has options for setting the message display.

- Choose Refresh to update the configuration information.
- Choose Options to change the polling time for the libraries.

Enable/Disable Library

In the SLC main window, select a library and click Enable or Disable, then confirm your choice in the dialog box that appears.

Note

You cannot disable the ODU or optical library in manual mode.

When an MSAR library is highlighted, the Enable/Disable buttons will be disabled or grayed out.

Enabling or disabling a storage library switches the status of the robotic arm between Normal mode and Disabled mode. When you disable a library, you are disabling the robotic arm, not the drive itself. The gripper moves to its home position and the system cannot send commands to the disabled library.

Calibrate Library

For FileNet OSAR storage libraries only, this function calibrates a library when the gripper is not properly aligned with the slots (such as after the library is moved or after an earthquake). Otherwise, your library should not require calibration. Calibration time varies from about 20 minutes to over an hour, depending upon the size of the library.

To calibrate a FileNet OSAR library:

- 1 On the SLC main window, select a library and click the Show button.
- 2 On the Library Configuration window, choose Calibrate Library from the Miscellaneous menu.

A warning message displays the time required to calibrate.

3 Click OK when ready to start calibrating, or click Cancel if you don't want to start the process.

Enable/Disable Drive

Use this function to enable or disable a drive.

- 1 In the Library Configuration window, select a drive.
- 2 Click the Enable or Disable button.
- 3 Click OK at the prompt to confirm your selection.

Enable/Disable Slot

Use this function to change the status of a slot. Your service representative directs you to use this function to disable a defective slot in the storage library. You can also use this function to re-enable slots that were previously disabled. When you disable slots containing media, the robotic arm moves the media to other available slots. If all slots are full, the media stays in the disabled slot and must be manually ejected.

Note

This feature does not apply to MSAR libraries.

- 1 On the SLC main window, select a library and click the Show button.
- On the resulting Library Configuration window, select Enable/Disable 2 Slot from the Miscellaneous menu.
- 3 Enter the number of the slot you want to enable or disable and the current status of that slot appears (enabled or disabled).
- 4 To change the slot's status, click the Enable or Disable button.

5 Click OK at the prompt to confirm your selection.

Enable/Disable Grippers

For FileNet OSARs only, this function displays the current status of the grippers and lets you change that status. Gripper error messages display in the Informational Messages area of the SLC main window. The system event log lists errors in the following format:

ARM (xx) Command: yyyyy Fault status: zzz

where:

xx is the ID of the process reporting the error yyyyy is the command sent to the storage library zzz is the fault status code sent by the optical unit

Write down this error message, then call your service representative for instructions on which gripper to disable. If that gripper is holding media, a service representative must move the media to the proper location before you can access it.

Note

This feature does not apply to MSAR libraries.

To disable a gripper:

- 1 On the SLC main window, select the appropriate library and click the Show button
- 2 On the Library Configuration window, choose Enable/Disable Grippers from the Miscellaneous menu.
- 3 Select the gripper your service representative tells you to disable.

You cannot disable both grippers at the same time. To enable a gripper, select Enable Both Grippers.

4 Click OK at the prompt to confirm your selection.

Media Management Functions

The media management functions described here are accessible from the SLC main window's buttons and menus or from the Library Configuration window's buttons and menus.

For information about moving MSAR surfaces or files or backing up MSAR media refer to **MSAR Procedures and Guidelines for Image** Services.

Preformat Media

Preformatting prevents storage media from being inserted incorrectly. We recommend that an experienced operator preformat all new storage media so media are available when needed. Preformatting marks the surfaces of storage media as Side A and Side B. SLC assigns surface IDs when the media are first written, not during preformatting.

Note

This feature does not apply to MSAR media.

Use the following procedure to preformat storage media used in an ODU. Do NOT use this procedure for storage libraries operating in normal mode.

To preformat media:

Select the ODU on the SLC main window and click the Show button. 1

- 2 If the drive you want to use contains any media, disable the drive and eject the media.
- 3 Choose Preformat from the Library Configuration window's Media menu and specify the drive you want to use.
- 4 When prompted, insert the storage media in the correct orientation.
- 5 When prompted, flip the media to the other side.

Insert Media

Use this procedure to insert either new or previously-used storage media into the library. Do not use this procedure to insert media from other systems. See "Incorporate Media" on page 378 and "Importing Documents" on page 374.

When inserting raw (unused) erasable media, the system must first pre-erase the media. This process can take up to 15 minutes per side. To keep a drive available for higher priority tasks such as reading for retrieval requests, the process is interrupted as necessary. Because of these interruptions, the wait can be much longer than 15 minutes before the system begins writing to the media.

Optical Storage Libraries

To insert storage media in a particular storage library:

- 1 Select the library on the SLC main window and click the Show button.
- 2 On the resulting Library Configuration window, select Insert Media from the Media menu.
- 3 Insert the media, as prompted.
- 4 Click OK to indicate you inserted the media (click Cancel if you do not insert any media).

When inserting raw erasable, preformatted media in a storage library, the system displays an informational message box, letting you know its progress:

Formatting started. This may take a while. Formatting completed.

These first two "Formatting" messages indicate the system is pre-erasing the media to zero out sectors, pre-written by the manufacturer.

5 If the media was already labeled, the identifying information appears in the blank fields.

Note

The system must keep one slot free for media currently in a drive. Therefore, if accepting media will not leave a free slot, the system accepts the new media then immediately ejects media (the media unused for the longest period of time) to free up a slot.

Optical Disk Units

Note

When inserting new media, you must insert it so the system writes on side A first. To assure that inexperienced operators do not ruin media by inserting it the wrong way, a knowledgeable operator should preformat blank media.

To insert storage media into an ODU:

- 1 Disable the drive in the Library Configuration window.
- 2 Eject any media in the drive by pushing the ODU's eject button.
- 3 Insert the new media.
- 4 Enable the drive in the Library Configuration window.

When inserting raw erasable, preformatted media in an ODU, the system displays an informational message box, letting you know its progress:

Formatting started. This may take a while.

Formatting completed.

These first two "Formatting" messages indicate the system is preerasing the media to zero out sectors, pre-written by the manufacturer. Following the two formatting message, an error message may occur:

Format operation failed

This message indicates the system attempted to label the media, but failed because the media had been inserted using the wrong orientation. When this happens, reinsert the media in the ODU using the proper orientation, then repeat the procedure.

MSAR Libraries

Note

When inserting new surfaces into a full MSAR, the system will prompt you to confirm this. If you answer yes, the system ejects the surface it determines is the least used surface. Be aware that all MSAR libraries on a server must be in Backup mode or all must be in Normal mode in order for you to successfully execute the insert media operation. If you only have some of the MSAR libraries on a server in Backup mode and you leave the rest in Normal mode, the insert media operation will not be allowed. For more information, see "Backup Library" on page 432.

For information about choosing the MSAR media to eject, see the MSAR Procedures and Guidelines for Image Services.

For information about inserting MSAR media, see **MSAR Procedures** and Guidelines for Image Services.

To insert media in a particular library:

- 1 Select the library on the SLC main window and click the Show button (or double-click on the library).
- 2 On the resulting Configuration for Library window, select Insert Media from the Media menu.
- 3 A dialog box entitled Select MSAR Surface Data File to Insert into Library displays. Enter the location of the MSAR surface file. MSAR data (.dat) or link (.lnk) files will be accepted in this field.

Note

Do not manually modify the contents of link files (using vi, for example) as you may get an error (such as <202,100,10> FCL error) when you

attempt to insert that MSAR surface. The editor adds extra characters when it is saved.

4 Click OK to indicate that you want to insert the media you have selected or click Cancel if you do not want to insert media.

Eject Media

Use this procedure to eject storage media based on surface ID or location. SLC completes pending writes before ejecting the media. For information about the logic involved in ejecting media, see "Appendix C – Logic for Retrieving Surfaces and Ejecting Media" on page 619.

Storage Libraries

When you know the storage media's location (otherwise, see "Eject Media by Surface ID" on page 446):

- Select the library on the SLC main window and click the Show button. 1
- 2 In the resulting Library Configuration window, select Eject Media from the Media menu.
- 3 In the dialog box that appears, click the radio button for your situation:
 - Click the Drive button for media in a drive.
 - Click the Slot button for media in a slot.

You can also enter the surface ID in this dialog.

4 In the box that appears to the right, enter the drive number, slot number, or surface ID.

- 5 Click OK.
- 6 Remove the media from the I/O station.

If your storage library is disabled, the system puts the request to eject media in a queue. When you restore normal mode, the system executes the request.

Optical Disk Units

To eject storage media from an ODU:

- 1 Disable the drive in the Library Configuration window.
- 2 Eject the media in the drive by pushing the ODU's eject button.

Eject Media by Surface ID

Use this procedure to move storage media out of a storage library (see "Optical Disk Units" on page 443 for ejecting media from an ODU). Before ejecting storage media, check the Detailed Surface Info report (see "Reports" on page 457) to be sure no pending requests exist for either side of the media.

From the SLC main window:

- 1 Select Eject Media (by surface ID) from the Media menu.
- 2 Enter the surface ID in the dialog box that appears.
- 3 Click OK or press Enter.
- 4 Remove the media from the I/O station.

From the Library Configuration window:

- 1 Select Eject Media from the Media menu.
- 2 In the dialog box that appears, click the radio button for your situation:
 - Click the Drive button for media in a drive and enter the drive number.
 - Click the Slot button for media in a slot and enter the slot number.
 - Click the Surface ID button and enter the surface ID.
- 3 Click OK or press Enter.
- Remove the media from the I/O station. 4

If the library is disabled, the system puts your request in a queue. When you restore normal mode, the system executes your request.


Enable/Disable Media

This option enables and disables automatic selection of a surface. given the family name. It does not physically disable I/O to a surface, and it does not interfere with existing read or write requests, since these operations access media by surface ID.


Talk to your service representative before disabling media. When you disable and remove partially full storage media, you have the option of disabling that media for reads or writes of either side. If you disable them for writes, you may need to know which media becomes the new current write surface. See "Media Family Information" on page 467.

To enable or disable media:

1 Choose Enable/Disable Media from the SLC window's Media menu. The Enable/Disable Media dialog box opens, along with a pop-up window, prompting for surface ID.



Enter one of the surface IDs and click the OK button. 2



Choose the button representing the surface you specified or both surfaces.

If necessary, you can change the surface ID shown at the top of the window.

3 Select the desired action at the bottom of the window.

The following table describes the result of a selected action in the Enable/Disable Media dialog.


Action	Result
Enable reads and writes (active)	Media enabled for active writes can be written by any write request queued to the family to which the media belongs. When you enable media for writes, you must enable Side B or both sides. You cannot enable just Side A. In contrast, you can enable either or both sides for reads.
Enable reads and writes (nonactive)	Nonactive media are written only if you are using clustering and you write to a cluster that already exists on the specified media.
Enable reads only	The surface is disabled for writes only.
	If you disable Side A for writes, the system switches any write requests to Side B. If you disable Side B or both sides for writes, the system switches write requests to the next write surface. If no next write surface exists for this family, the system asks for new media.
Disable reads and writes	The surface is disabled for both reads and writes. The system satisfies retrieval requests by attempting to read from the transaction log copy of the document. If the transaction log is not available, you must enable the media to complete the retrieval.
	Disabling a source surface for both reads and writes during a media copy has no effect on the copy operation.
	Disabling media containing documents that are being imported has no effect on the import operation.

Change Media Type

You cannot change the media's type once storage media has been inserted into the storage library and been written to.

Before inserting transaction log media into a library:

On the SLC main window, select Change Media Type from the Media 1 menu to display the following dialog box.



2 Enter the surface ID of the media to be changed.

The surface ID you enter must be the newly-assigned surface ID for the importing system. You can get the surface ID information from the Local/Foreign IDs report (in the SLC main window's Media menu).

- 3 Click the Type button to see the current type.
- 4 Select the appropriate media type from the pulldown list.
- 5 Click OK to change the media type or Cancel to close the dialog without changing anything.


Change Media Family Name

Use this function to change the family name associated with primary storage media (for example, if you import to the wrong family or add existing media to a new family). The family name must already exist. You cannot change primary families to transaction log families or change transaction log families to primary families. This function removes the storage media from the list of current surfaces for the old family. It does not automatically put the media in the list of currently writable surfaces for the new family. To write to the media via the new family, you must enable the media to explicitly start writes to its surfaces.

Before you can change the media family name, the following must be true:

- Neither media surface is open.
- The surface is not in any library.
- Neither media surface has any requests pending.
- Neither media surface is in use by any process.

1 Select Change Family Name from the SLC main window's Families menu.



- Enter the surface ID of the storage media you want to change, then 2 click the Family button.
- 3 Select the new family name, then click OK to change the name, or Cancel to close the window without changing anything.

Each optical disk has at least one document header file. This file contains a record of the documents committed over a period of time. The system creates the first file; you create additional files as needed. When you create a new document header file, the file contains a record of the documents committed from that point on, until you create the next document header file.

Note

The Create Document Header File operation is not allowed if an MSAR library is already in Backup mode.

If you are entering and committing documents on one system for export to another system, use document header files to group documents chronologically. You can then take the media to another system and import only the documents listed in the specified document header file.


The name of a document header file looks like this:

FN DESCRIPTOR 19980923 142509

The name always starts with FN DESCRIPTOR. The rest of the name consists of the date and time (yyyymmdd hhmmss). Time is in 24-hour format.

To create a document header file:

1 Select Create Doc Header File from the SLC main window's Media menu.



2 Enter the ID of the surface you are writing.

> If you are writing the ID on imported (foreign) media, use the new surface ID that was assigned by your system during the import.

3 Click the Generate button.

> The program tells you that it is creating the file and then displays the name for a short while. (The import function can display the names of all document header files.)

Identify Media in Library

This option creates a new map of where media are located by reading the media. Make sure at least one drive is enabled before selecting this option. If all drives are disabled, the system cannot read identification information on media, but can only mark slots as containing unidentified media and mark surfaces as lost.

The storage library maintains an internal map that keeps track of the media in each slot. If the map information does not match the actual location of any of the media, use this function to tell the library to create a new map. You might use this function when the system cannot find media that you know is in the library.

The system checks each slot for media, then inserts media it finds into a drive to read and display the media's volume (identification) information. Disabled slots containing media are not identified.

#New material


The storage library must be in normal mode to use this function.

CAUTION

Because this process reads all media in the library, it can take several minutes to an hour to complete.

To identify the media in a library:

- On the SLC main window, show the appropriate library's information. 1
- On the Library Configuration window, choose Identify Media in Library 2 from the Miscellaneous menu.




Click OK to acknowledge the time warning and start the process. A 3 message box lets you know the process has been queued.

Reports

The SLC main window and Library Configuration window contain menu selections that provide reports on slots and drives, document committal information, media surfaces, how space is used on the media surfaces, and media family information.

Slot/Drive Map


Select Slot/Drive Map from the Library Configuration window's Reports menu to see which surfaces are in each slot or drive and whether a drive or slot is disabled. An asterisk indicates a disabled slot.



Local Statistics

This report lists the number of pages retrieved and the number of documents committed and deleted for any day. These statistics are kept for your local system. To access to these statistics, ask your service representative about configuring this option. It requires extra system resources.

To view the local statistics, select Local from the SLC window's Statistics menu. If you turned quick logging on since you last recycled the Image Services system, the Date field displays the current date, as shown in the following example.



If quick logging is not running when you select this option, the Date field shows the date format mask, using the default date mask configured on your system platform. The system also displays the following error message:

Quick Logging is currently not running. Local statistics will still allow you to view previously logged data.

You may enter the date of the local statistics log file, using the format specified in the date mask.


Click the Refresh button to generate current statistics. To print the report, click the Print button. (For details on using the print dialog box, see "Print File or Report Dialog Box" on page 543.

Remote Committals



This report lists the number of pending remote committals. If page cache seems fuller than expected, check this report for pending remote committals. Then you can contact the remote system to see when committals can resume.

To view the remote committals report, select Remote from the SLC window's Statistics menu.




Pending Surface Requests

This report contains information about all media with outstanding requests even if the media are not in any library. If no media have requests, the following message displays:

There are no known media with pending requests.

Choose Pending Surface Requests from the SLC window's Media menu to display a report similar to the following.




Use the horizontal and vertical scroll bars to view all information in the report.

Detailed Surface Information

Before you remove media, always check this report to make sure that there are no pending requests for either side of the media

Choose Detailed Surface Info from the SLC window's Media menu. enter the surface ID number, and click on Display.



Specify Surface

To view surface information and statistics, enter the ID number of the surface and click the Display button.

You can also print this report using the Print option from the File menu. (For details on using the print dialog box, see "Print File or Report Dialog Box" on page 543.)

Surface Description

This box describes the specified surface. The system serial number applies to the foreign system that originally wrote the media. The Sector Size field provides the media size.

MSAR-specific information includes the following:

The **MSAR Read Only** field states whether or not the user can only read this surface. When an MSAR surface file encounters the **Out of** space error, the MSAR Read Only flag will be set to TRUE. This is the only condition that will cause the MSAR Read Only flag to be set. During the time the flag is set to TRUE, only the read requests for the surface can be processed.

The **MSAR FileName** field is a scrollable display window which indicates the last known location of this surface.

Surface Statistics

This box summarizes surface activity, such as the number of active and deleted documents and sector usage.

Local/Foreign IDs

This report shows you the system serial number (ssn) of the system that wrote the media and the media's current surface ID. If written by a different system, the original surface numbers appear along with the serial number of that system.

Choose Local/Foreign IDs from the SLC window's Media menu to display the Local/Foreign ID report.


The local surface ID is a unique number that identifies a surface on your current (local) system. No two surfaces on the local system can have the same surface identifier.

The foreign surface ID displays only if the surface was written by a different (foreign) system. If this is the case, the system serial number for this media is different from your local system serial number. Generally, systems are configured so surface IDs do not overlap.

To see information for a particular surface, click the Go To button and enter the surface ID. Indicate whether this is the local surface ID or a foreign ID.

Media Surface Summary

Select Media/Surface Summary from the Library Configuration window's Reports menu to display a report of the pending read and write requests for all surfaces, regardless of whether the media are in a library.




An asterisk (*) indicates which surface is currently accessible by the drive. The report does not include reads and writes that have taken place for any media that are presently outside of the storage library.

The table below describes the contents of the columns of the media surface summary report.

Column	Description
Surface	Media surface ID
Location	Media location
High Reads	User retrieval requests
Medium Reads	Page read-ahead requests (retrieval software prefetches the page most likely to be displayed next)
Writes	Write requests
Low Reads	WorkFlo prefetches
Background	Background jobs, such as importing and copying media

Media Space Use

Choose Media Space Usage from the Library Configuration window's Reports menu to display a report of how space is used on each surface, including active and deleted documents.




This report also shows the percent of space available, average number of media sectors per document, average number of pages per document, and the estimated active sectors. The estimate of active sectors is a product of the two averages (number of active documents times average sectors per document).

Media incorporated from foreign systems do not have valid space use information until the import operation is complete for all documents on the media. The number of active and deleted documents on the media is only an estimate. Use these numbers primarily as an aid to determine when to consolidate media.

Media Family Information

The Media Family Information report lists information about a media family.

Select Media Family Information from the SLC window's Families menu. Enter the name of the media family (or select a family name from the pulldown list). A report similar to the following displays.



This report lists the ID of the current write surface for the given media family and lists up to eight future surfaces. The future surfaces are reenabled surfaces, available for future writing. When the system finishes writing the current surface, it looks first in the future surfaces list and uses this older media before asking for new media.

Commands

This chapter contains general information about command syntax, command options, and using scripts, followed by an alphabetic reference to many of the commands and utilities provided by your Image Services system.

UNIX commands such as **history** and **more** are not included in this chapter. See your operating system manuals for information about using these and other standard operating system commands.

Logging On to the Security System

Some of the tools described here require you to be logged onto FileNet security. A command that requires security may let you display help text without logging onto security.

Windows Server Users Logon

Log onto FileNet, then enter the commands in this chapter in a command prompt window.

UNIX Users Logon

To log on to FileNet security from the command line, enter:

fnlogon

The fnlogon command gives the current command line shell security access. Therefore, if you switch to another shell, you need to run fnlogon from the new shell.

Use Control+d to exit fnlogon at any time.

CAUTION

Image Services Toolkit (ISTK) programs must explicitly log on to FileNet security as they do not inherit the logon information from the fnlogon program.

The fnlogon command prompts for information:

```
costa3(kehr)/home/kehr> fnlogon
FileNet user name: SysAdmin
FileNet password:
FileNet security service (CR = local service):
Program (CR = default shell):
   Name : SysAdmin:costa3:FileNet
   Logon # : 274
   Last successful logon:
      time: Tue Jan 16 14:15:43 2003
      where: WS001@135.0/2.6:costa5:FileNet
   Last unsuccessful logon:
      time: Fri Jan 12 14:15:29 2003
      where: WS001@135.0/2.6:costa5:FileNet
      error: < 92, 2, 2 >
fnlogon: executing /bin/csh...
```

After entering your FileNet user ID and password, press Enter to access the local system, or enter the name of the security service for a networked system:

<SecurityService>:<domain>:FileNet

CAUTION

When you log on to a remote security service, you may encounter errors if subsequent applications you launch are not designed for remote operation.

In response to the Program prompt, choose one of these actions:

- Enter the name of the command you want to run (for example, CSM_exim). When you quit that command, fnlogon also quits.
- Press Enter to go into the shell prompt and type the name of the command you want to run at the command line. When you guit that command, fnlogon continues to run so you can access other commands without logging on again.

GUI Commands

Some commands were formerly issued from the command line but now have a graphical user interface (GUI).

initfnsw

The **initfnsw** command, which stops and starts the FileNet software without rebooting, is accessed from the Task Manager.

See "Software Control Buttons" on page 321 for information about using the Task Manager to start, stop, restart, or prepare the FileNet environment for backup or restore operations.

νl

To view all current event log files, use the System Monitor's Event Log instead of the vI tool (see "Monitoring Event Logs" on page 327).

You can also use the **less** utility (UNIX platforms only) to read a single event log (see "less" on page 509).

whatsup

To display a list of currently-running FileNet programs loaded into memory of the station you are logged onto, use the Task Manager instead of the whatsup command (see "Check the Active Processes" on page 330).

Commands Overview

This chapter provides the following information for each command:

- command name
- brief description of what the command does
- syntax of the command

A command that you type at the command line can have three parts: command name, arguments, and variables.

An argument enclosed in square brackets is optional (do not type the brackets). Variable information that you supply is enclosed in angle brackets (<variable>); again, do not type the brackets. Otherwise, type the command as follows:

command [arg1] [arg2]...argn] <varname>

Most commands are terse. Arguments are generally one or more letters preceded by a minus sign. The examples provided use the ddexim command (see "ddexim" on page 488).

Arguments

Arguments modify commands. For example, the ddexim command exports or imports data dictionary information. To export the data dictionary and display the results on the console, enter:

ddexim -e

You can combine arguments in the same command. For example, to export everything in the data dictionary except the media families, apply both the **-e** and the **-sf** options:

ddexim -e -sf

Redirecting Output to a File

By default, the output of many commands is directed to the device named stdout (standard out), which means it displays output of the command on the console. You can save the output to a file by using the redirect (>) symbol followed by the receiving filename:

command > filename

For example, to create a file containing the contents of the data dictionary, use this command:

ddexim > dictionary

If the file does not exist, the redirect command creates it. If the file exists, the command overwrites it. To append the output of a command to an existing file, use the >> symbol:

ddexim >> dictionary

CSM exim

The **CSM_exim** command exports (backs up) objects from a specified logical cache to tape or to a disk file. CSM exim also has an import function, though we recommend that you request help from your service representative to import objects into your production system's cache. This manual does not include information about the import function.

If you run CSM exim while cache is changing, the tape may not be synchronized with the contents of cache at the end of the backup. CSM exim does not export or import batch or write request information from the transient database.

CSM exim requires 10 bytes of memory for each object backed up. If you back up a very large number of cache objects, CSM exim may fail with an out-of-memory error during the backup. If you get an out-ofmemory error, use the Cache Export/Import Program (accessed from the Application Executive's Applications menu), which does not have this memory constraint.

Note

The Cache Export/Import Program uses the same tape and file formats as does CSM_exim, allowing data interchange between the programs. Cache Export/Import does not back up or restore individual objects, read an ASCII file to specify objects, or restore cache data to a different cache. See the "Backup" chapter in your System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server for details on the Cache Export/Import Program. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

To help you decide which program to use for backing up cache, see "Cache Backup" on page 54.

CSM_exim tape request messages display in the console window. When using CSM_exim to back up cache objects to tape, you must run the command in the foreground on the server to which your tape is physically connected.

The CSM exim command:

- Finds the objects in cache
- Creates the directory including the cache ID, system serial number, object ID, and page
- Displays the directory unless the -q option was specified
- Writes the objects to tape or a disk file
- Creates the report file. To examine the report file, use a text editor (more, less, vi, etc.).

Determining What is Backed Up

As soon as you start the command, CSM exim begins compiling a list of the objects to be backed up. It cannot back up objects written to cache after or while compiling the list. The command output includes a list of backed up objects, so you can see which objects are on the tape.

Compressed Data

The Image Services software compresses images stored in cache. If you are using a tape drive that compresses data, the tape drive cannot further compress the cached images.

Running CSM_exim

To run CSM exim:

- UNIX users, log on to FileNet security using fnlogon or as a member of the fnadmin group and enter the CSM_exim command on the command line.
- Windows Server users, log onto FileNet using the Application Executive and enter the CSM exim command in a command prompt window.

The two most common forms of the CSM exim command are:

```
CSM exim -e -d tape -c <cachename>
   or
CSM exim -e -d <filename> -c <cachename>
```

The following table describes each option you can use with the CSM exim command.

CSM_exim options

Option	Description
-е	Exports cache objects to tape or file
-d <filename></filename>	Specifies the destination filename. If you omit –d , the filename defaults to CSM_EI_DAT. If you export to magnetic disk, you can specify a path and filename of up to 75 characters, but the filename cannot exceed 10 characters.
	CSM_exim -e -r -d /fnsw/local/tmp/CSMexp
	Enter the word tape to specify a local drive (if one exists) or enter the three-part NCH name of a tape drive. The CSM_exim default tape drive is the one installed at the server where you enter this command. If this server has no tape drive, then CSM_exim uses the drive configured as the system default.
-p <pc filename=""></pc>	Specifies a PC partition and destination filename. A colon (:) is required when you specify the drive letter as part of the pathname. For example:
	CSM_exim -e -r -p D:\fnsw\local\tmp\CSMexp
-c <cache name=""></cache>	Each occurrence specifies a logical cache. For example:
[-c <cache name="">]</cache>	CSM_exim -e -c folder_cache1 -c fillin_cache1
-n	Displays selections, such as cache and object IDs, from the specified cache, tape, or file. However, no export occurs. Used for information only.
-o <objects></objects>	Specifies individual objects for export/import. Used with the -c option, specifies individual objects within specific logical caches.
-v	Displays the location of cache objects (cache:domain:organization) during the export process
-q	Turns off display of cache object attributes (cache ID, system serial number, object ID, etc.) during export
-1	Exports only locked objects. Use –I (lowercase L) to back up only uncommitted documents residing in a cache with ageable documents
- \$	Suppresses excluded cache objects from displaying on the export list. You can excluded objects using the –I or –t option.

CSM_exim options, Continued

Option	Description
-t <time></time>	Export only the objects with either the creation date time or the "lastupdate" date time later than the specified time. The time is in the format:
	"mm/dd/yyyy hh:mm:ss am pm"
	Be sure to include the quotes and either am or pm .
-f <filename></filename>	As an alternative to specifying the -c option multiple times, the -f option specifies a file containing a list of logical caches in this format:
	cache
	<cachename1>[,<cachename2>]</cachename2></cachename1>
	or
	cache <cachename1> objects <object1>[,<object2>] [<cachename2>] objects <object1>[,<object2>]</object2></object1></cachename2></object2></object1></cachename1>
	To specify a cache name, enter the name (for example, bes_cache1) on the local system or the three-part NCH name for a remote system.
	To specify objects, include the system serial number, document ID, and page number, or use an asterisk (*) to indicate all items. For example:
	<ssn> <document id=""> *</document></ssn>
	indicates all pages of one document for this system.

Output Results

CSM exim produces three types of files in the directory from which you run the command:

- CSM EI RPT.cess ID> is an ASCII file containing the same information that displays on the console window when running in the foreground.
- The other two files default to CSM EI DAT.<cache ID> and CSM EI DAT.DIR. These are binary files that list what was exported. If you supply a filename for the -d or -p option, the filename replaces the CSM EI DAT portion of the name.

Note

When backing up to tape, the system writes the CMS_EI_DAT.<cache ID> and CSM EI DAT.DIR files directly to tape. One CSM exim can produce multiple .DAT files; it produces one .DAT file for each backed up cache.

CSM tool

To run **CSM** tool, which includes several options for monitoring cache, you must first log on to FileNet security. To start CSM tool, enter the following at the server where you want to check the cache:

CSM tool

The following appears on the screen:

Type '?' for help

<CSM tool>

Entering a question mark displays the following help text:

```
<CSM tool>?
Upper case letters are minimum abbreviations.
'CHeckcache' - Checks the caches for overlaps or missing free space.
CHECKSum <cache id> <ssn> <object id> <page>
Displays checksum saved in database and checksum computed on data
'DUmpobject <cache_id> <ssn> <object_id> <page>' - Dumps object in hex.
'HArdcopy [<filename>]' - Puts copy of all output in file <filename>
        if present, or turns off hardcopy if <filename> not present.
'Listobjects [<cache_id> [<ssn> [<object_id> [<page>]]]] [desc]' - Lists
        objects from starting location in ascending or descending order.
'OBjecttofile <cache_id> <ssn> <object_id> <page> <filename>' - Writes
        the object indicated to a UNIX file with name '<filename>'.
'Statistics [<cache_id>] [{long | sect}]' - Prints statistics for a single
        cache or all caches, in short (default), long, or sector format.
'Ouit' - Exits program
```

Use the cache ID (a number) instead of the cache name when specifying a single cache. Use the **listobjects** command to find the cache ID number.

checkcache

The **checkcache** command checks the caches for overlaps or missing free space or for objects that appear to reside on a nonexistent platform.

CAUTION

If you execute checkcache while the FileNet system is running, all system access is denied to all caches on the server until checkcache completes.

CAUTION

Checkcache performed on a large cache could take hours to complete.

The system displays an error if it detects any problems in the cache. If it finds no problems, the system displays a message similar to this:

<CSM tool>CH Building file /fnsw/local/tmp/cache.list... Sorting into file /fnsw/local/tmp/cache.sort... Check complete, 0 inconsistencies found.

listobjects

The listobjects command displays a list or range of objects with their corresponding attributes (cache_id, ssn, object_id, page, and max_ length) in ascending or descending order. In the following example output from the listobjects command, some objects in page cache are not marked as ageable. These objects are locked.

<csm_too1>1</csm_too1>					
cache_id	ssn	object_id	page	max_length	
1	10291	20026080	9	232	aqeable
1	10291	20026080	1	359889	ageable
1	10291	20026080	2	359889	ageable
1	10291	20026080	3	359889	ageable
1	10291	20026080	4	359889	ageable
1	10291	20026080	5	359889	ageable
1	10291	20026080	6	359889	ageable
1	10291	20026080	7	359889	ageable
1	10291	20026080	8	359889	ageable
1	10291	20026080	9	359889	ageable
1	10291	20037321	9	116	-
1	10291	20037321	1	6562	
1	10291	20037323	9	124	
1	10291	20037323	1	3064	
1	10291	20037462	1	3072	
1	10291	20037465	1	3072	
1	10291	20037466	1	3072	
1	10291	20037467	1	3072	
1	10291	20037468	1	3072	
1	10291	20037469	1	3072	
1	10291	20037470	1	3072	
'CR' = line,			paging		q' = exi

The following table lists and describes the action of each option you can use with the listobjects command.

Option	Description
[<objectid>]</objectid>	Lists the specified objects
[<objectid range="">]</objectid>	Lists the specified object range
[FOR <count>]</count>	Lists a specified number of objects in ascending order
[FOR <count>] [desc]</count>	Lists a specified number of objects in descending order
LF or LB	Lists objects forward (LF) or backward (LB) If you do not specify <objectid>, the system starts listing from the most recently displayed objects. If it hasn't displayed any objects, the system starts listing with the first or last object.</objectid>

You can repeat all listobjects commands by pressing the Return key. When you use the Return key to repeat LF and LB, the system lists the next group of objects. The group size is determined by the <count>; if you do not specify <count>, the system continues listing to the end of the objects in cache.

You can get detailed information about each option in the CSM tool online help.

Statistics

The statistics command (S) displays statistics for a specified cache or for all caches if a cache_id is not specified. You can specify short, sector, or long format for your output listing. If you do not specify an output format, a short report is the default. Each line displayed in the table provides information about one type of cache.

Short Format

To see the short form of statistics for all caches, enter **S**.

<csm_tool>S</csm_tool>				
Cache Id Name		% locked	% full	% free
1	page_cache1:moorea:FileNet	7	12	88
3	bes_cache1:moorea:FileNet	24	24	76
4	sys_print_cachel:moorea:FileNet	0	0	100
5	app_print_cache1:moorea:FileNet	1	1	99
*	Physical space summary	16	17	83
<csm_tool< td=""><td>></td><td></td><td></td><td></td></csm_tool<>	>			

Sector Format

Entering the sector option (S sect) displays sector format, showing a summary of the number of 1K sectors.

```
<CSM tool>S sect
Cache Id Name
                                            # locked # full # free
_____
                                            _____
                                                            -----
      1 page_cache1:moorea:FileNet
                                                1417
                                                       2395 18904
      3 bes cache1:moorea:FileNet
                                               14901 14901 48996
      4 sys_print_cache1:moorea:FileNet
                                                  0
                                                          0
                                                             21299
       app_print_cache1:moorea:FileNet
                                                            31946
      * Physical space summary
                                                      17298
                                               16320
                                                             89197
<CSM tool>
```

Long Format

The long form for cache statistics (**S long**) shows this information:

```
<CSM_tool>S long
Statistics for cache #1, name = 'page_cache1:moorea:FileNet'
min_cache_sectors 21299
                          locked_sectors 1417
                                                 inuse_sectors 2395
max_cache_sectors 21299 locked_objects 17
                                                 inuse_objects 65
free_sectors
                 18904
                         ageable
                                                 refents
self_cleaning
Statistics for cache #3, name = 'bes_cache1:moorea:FileNet'
min_cache_sectors 10649
                          locked_sectors 14901 inuse_sectors 14901
max_cache_sectors 63897
                         locked_objects 153
                                                 inuse_objects 153
free_sectors
                 48996
                          ageable
                                        F
                                                 refents
self cleaning
Statistics for cache #4, name = 'sys_print_cache1:moorea:FileNet'
min_cache_sectors 10649 locked_sectors 0
                                                 inuse sectors
                                                               0
max_cache_sectors 21299
                         locked_objects 0
                                                 inuse_objects
```

free_sectors self_cleaning	21299 F	ageable	F	refcnts	F
Statistics for cac min_cache_sectors max_cache_sectors free_sectors self_cleaning	10649 31948 31946	me = 'app_print_ locked_sectors locked_objects ageable	2 1	inuse_sectors	2 1 F
Physical space sum reserved_sectors max_cache_sectors free_sectors	53246	locked_sectors locked_objects largest_fs_blk	171	inuse_sectors inuse_objects	17298 219
Prefetch duration					

cti

The **cti** (count index values) tool counts the number of unique non-null values recorded in the index database for a specified index.

cti <index name>

The information returned by cti is useful for planning retrieval strategies. In general, the more unique values in the database, the fewer potential matches exist (if the distribution is fairly even). To choose the most appropriate index to use as a primary key, choose the index with the most unique values.

You can also use cti to acquire statistics about information in your database. For example, if you have an account number index, you can use cti to find out how many accounts you have. Since this tool works with system indexes as well, you could also use cti to count the number of documents in your system by specifying **F_DocNumber** as the index name.

To use cti, enter the command at the index server. When entering the index name, spell and capitalize the name exactly as it exists in the database. The index must be a retrieval key but does not need to be unique (F DocNumber is usually the only unique index on the system). Since cti gives you the number of unique values, an index with duplicate values is counted only once. You can run cti during normal processing, but the information returned will be current only for batches that are already committed.

Here is a sample report for a very small database:

```
costa5(kehr)/home/kehr> cti F DOCNUMBER
Direct database count is used. There will be no status
prompts.
Number of distinct values for the index is: 715
Total number of index values examined is
```

ddexim

The **ddexim** command exports or imports information about indexes, document classes, media families, forms, menus, clusters, autoindexing (labeled aperture card indexloc in the output), WorkFlo queues, and storage media surfaces (if specified).

This information is collectively called the data dictionary. Along with backing up this information, ddexim is useful for setting up remote systems with identical parameters.

CAUTION

Do not import a data dictionary into an active system. See your service representative if you need to perform other operations.

To run ddexim, you must be logged in to FileNet security services. Displaying the help text (by typing ddexim with no parameters) does not require security services logon.

Use the export option to write the data dictionary to an ASCII file:

ddexim -e > <filename>

Use the import option to copy the data dictionary information to a new entry system:

ddexim -i <filename>

To cancel ddexim:

- If running in the foreground, press Control+c.
- If running in the background (UNIX only), enter:

$$kill -31 < pid>$$

where <pid> is the process ID of the ddexim program.

The following table describes each ddexim option.

Option	Action		
–е [<filename>]</filename>	Exports the data dictionary and displays the results or writes it to a file when you specify the filename		
-i <filename></filename>	Imports information from <filename></filename>		
-s <x></x>	Use -s with o, d, f, or w to skip import or export of specified information:		
	-so Skips RDBMS information (document classes, indexes, keys, clusters, and menus)		
	-sd Skips dumped information (forms, autoindexing information)		
	-sf Skips media families		
	-sw Skips WorkFlo queues		
-sd -i <filename></filename>	Includes security definitions from <filename>. You may use this option to import security definitions assigned to document classes, as defined in the specified file. Using this command on an Windows Server with an MSSQL Server installation requires special procedures.</filename>		
-c <docclass></docclass>	For a specified document class, exports or imports related indexes and families, but not surfaces. Repeat for each document class as needed.		
-as	Adds export or import of surface information (imports only if source and target system serial numbers are the same). You cannot specify –as when using –c .		
-m	During import, modifies preexisting structures, except for indexes and surfaces		

Option	Action
-v1	Writes the output at the IMS Release 3.0 version level*
-v2	Writes the output at the IMS Release 3.1 or later version level*

*Use this option to make the data dictionary compatible with an earlier release.

The following sample output shows several blocks of information in an ASCII file created by the -e option. The blocks include a document class, WorkFlo queue, family, and index.

```
class dialog {
description "dialog"
security_read_name "(ANYONE)"
security_write_name "(ANYONE)"
security_exe_name "(ANYONE)"
retent_disp delete
retent_base rel_to_entry
retent offset 12
pages_per_doc 0
tab_out_flag false
verify_images false
verify indexes false
batch total false
batch size 5
family name HPrint
no_catalog_flag false
index Description: required=f batch_totals=f verify=f
source=1
index TestDate: required=f batch_totals=f verify=f source=1
index TestNotes: required=f batch_totals=f verify=f source=1
index Description2: required=f batch_totals=f verify=f
source=1
index numeric: required=f batch_totals=f verify=f source=1
index processdate: required=f batch_totals=f verify=f
source=1
```

```
wflq a1 joec {
   nch object name WflServer
    desc read sec name "(ANYONE)"
    desc write sec name "(ANYONE)"
    desc exe sec name "(ANYONE)"
    content read sec name "(ANYONE)"
    content write sec name "(ANYONE)"
    content exe sec name "(ANYONE)"
    description "queue for testing dump operations."
    field: fld1 typ=1 len=16 unique=0 reg=t rendev=f disp=t
family tranlog {
    is primary false
    tran families
    interleave cnt 1
   disk type 5
    server 2: desired cur surfs=0 preferred osars=*
index Description {
    description "Description Uppercase Convert"
    internal name A31
    tvpe ascii
   maxstrlen 60
upper false
index TestDate {
    description "Date of the Test"
    internal name A32
    type date
    upper false
```

deldocs

The **deldocs** command deletes the specified document records from the Image Services index and MKF permanent databases. It can also delete selected documents from cache, given selected settings.

CAUTION

This command deletes all references to the documents from the Image Services index and MKF permanent databases. You can, however, reimport these deleted documents from storage media.

For details on deleting documents from cache, see <u>"Documents in Cache" on page 493</u>.

Other document deletion options include:

- Delete Doc/Folder option in Database Maintenance
- WorkFlo's delete document call
- Desktop application Query program's delete document option

You must be logged on to FileNet security (Windows Server users, logon through the Application Executive) before using deldocs.

Note

The deldocs tool will not delete a document if it is stored on a Single Dcoument Storage (SDS) device, if the SDS delete option is turned on for the SDS unit, and if the document retention setting is set to one of the following: infinite, Event Based Retention (EBR), or chronological retention. In the case of chronological retention, deldocs will not delete a document only if the date has not yet been reached. To ignore the SDS retention setting, use the -i option. For complete information about the deldocs tool, see the *Image Services System Tools Reference Manual*. To download this manual from the IBM support Web site, see "Accessing IBM FileNet Documentation" on page 19.

Document File List

To delete a set of documents listed in a file, enter the following command at the prompt:

deldocs -f <filename>

where <filename> is the name of a text file containing one document ID per line.

To create the file, you can use a WorkFlo program that identifies documents with the appropriate criteria (for example, older than a specified entry date, closed, and so forth).

Documents in Cache

The deldocs -f command behaves differently depending on two things:

- Whether the class assigned to the document is set to migrate to optical disk, and
- Whether the operator overrides this document class setting by modifying the Capture Committal Component.

Document Class Setting

When scanned into the FileNet system, the document is assigned to a class. The document then inherits any parameter values set for the class to which it is assigned.

The System Administrator configures classes using the Database Maintenance application, available through the Application Executive (Xapex). The Database Maintenance Classes option, Define/Update

Document Classes, enables the administrator to set the Migration to O.D. (Optical Disk) field to either Yes or No:

- Yes specifies that any documents associated with this class will be migrated to optical disk. It also flags the document as "ageable," unlocking it in cache. As long as the Capture Professional application does not override this setting before committal, deldocs does not delete the document from cache. Instead, either CSM daemon can remove the document from cache at an appropriate time or you may remove the document manually using CSM tool.
- No specifies that any documents associated with this class will not be migrated to optical disk. After committal has completed successfully, the document resides in cache as a "locked" document. As long as the Capture application does not override this setting before committal, the deldocs command deletes a locked document from cache.

Note

Use the CSM tool utility to see if a document in page cache contains a "locked" or "ageable" status.

Capture Committal Component

Under normal conditions, the Capture Committal Component inherits the default setting for the given document class, whether Migration To O.D. is set to Yes or No. If set to Yes, Capture shows Commit with Migration checked. If set to No, Capture shows Commit without Migration checked. If the operator modifies the Capture Committal Component, it overrides the setting assigned to the document class.

Erasable Media Surface

You can also use the deldocs tool to delete all the documents on an entire erasable media surface, then use the consolidation procedure in Background Job Control to erase the media without the time-consuming copying usually involved in consolidation.

To erase a surface:

- 1 Log onto FileNet security.
- At the command line, enter: 2

deldocs -s <surface#> -b

In place of <surface#>, enter the surface number of the media from which you want to delete documents. You must enter a separate command for each surface.

3 Perform the Consolidate Media function in Background Job Control (see "Consolidating Media" on page 398 for details).

Be sure to check the Erase Media box before beginning the consolidation. The consolidate process first determines that no documents exist to consolidate, then erases the source media.

- 4 Eject the media by surface number.
- 5 Reinsert media into library.

enlarge_ncol



The Microsoft® SQL Server™ RDBMS and IBM DB2 RDBMS requires you to specify a fixed length format in a user mask when you create a numeric index. You must specify the number of digits, called precision and scale, before and after the decimal point. If you try to increase the mask in Database Maintenance to an amount larger than the allocated amount, you get an error message.

CAUTION

The FileNet Image Services software **must** be shutdown when running enlarge_col.

The enlarge_ncol command increases the precision and scale of an existing numeric column in a Microsoft SQL Server database. It automatically adds padding digits to allow for future increases in the number of digits that make up the numeric index without creating another column.

This command performs the following sequential operations to accomplish the column enlargement:

- Renames the old column
- Creates a new column with a larger format
- Copies the data from the old column to the new column
- Inserts null characters in the old column to recover the space

Before this command enlarges the column, it displays a summary of the proposed updates with a confirmation prompt. You can enter \mathbf{y} to continue or \mathbf{n} to terminate the operation.

Limit the number of numeric index enlargements. Every index that is enlarged uses another column in the DOCTABA table of the index database. The space used by the old column is wasted.

Running the enlarge_ncol command

To run this command, follow these steps:

- 1 Shut down Image Services software.
- Start the relational database management system. 2
- 3 Determine the current precision and/or scale of the user index to be enlarged.

You can use ISQL to query the database for these values.

Enter enlarge ncol at the command line. 4

> If you do not specify options when you enter the command, the program prompts you for options. (See the options listed in the table below.)

5 When the command completes, restart Image Services software.

enlarge_ncol options

Option	Description
-i <indexname></indexname>	Specifies the numeric index to enlarge
-p <pre>-p <pre>cision></pre></pre>	Sets a new precision.
	Precision is the total number of digits on both the left and right side of the decimal point (excluding the decimal point, commas, etc.). The new precision value must be greater than the old precision value. The maximum value for Microsoft SQL Server is 30. The maximum value for DB2 is 30.
-s <scale></scale>	Sets a new scale. The new scale value must be greater than the old scale value. The scale value must be less than or equal to the precision value.
-r <rows></rows>	Specifies the maximum number of rows to update before committing (default: 2000)
-n	Sets old column data to null. If you do not specify this option, the old column contains the old data.
-h	Prints syntax help text

Tip

Use ISQL queries to determine the current precision and scale settings. Also, the enlarge_ncol command displays the current settings before prompting you for confirmation to update, as shown below:

Current database column is type numeric(21,7) New database column will be type numeric(22,8) Update (y/n):

Example

The following example is the result of a successful attempt to enlarge a numeric column (responses to prompts appear in bold type):

> enlarge_ncol Enter name of the numeric user index to enlarge: User_ index num Enter precision: 22 Enter scale (0-22): 8 User index 'User_index_num' will be updated as follows: Current database column is type numeric(21,7) New database column will be type numeric(22,8) Update (y/n): **y** Please wait... Copying data from column 'a38' to column 'a39' 96/08/12 17:21:23 22100 rows updated Successfully enlarged numeric index 'User_index_num'

fn_msg

The FileNet **fn_msg** command interprets FileNet error tuples.

fn_msg <error tuple>

The fn_msg command provides explanatory text for error numbers found in FileNet error logs. Enter the error tuple in either of the following ways:

A three-part value:

<nnn,nnn,nnn>

For example:

fn_msg 170,0,17

An 8-digit hexadecimal number:

<0xnnnnnn>

For example:

fn_msg 0xef00ab

gcp



Global file copy (**gcp**) is available on UNIX platforms.



gcp [-options] <sourcefile> <targetfile>

Use gcp to transfer files between two machines on the same network. For example, you can use gcp to transfer a file between two Image Services servers:

gcp OSAR:/tmp/logs/1/elog960414 /tmp/Flog960414

Specify a file or directory name on the remote machine:

<rhost:/path/file>

If you do not specify a path for a remote file or directory, gcp places the file in /tmp, with this exception: if the REMOTEDIR variable is set, gcp uses the path specified by this variable.

You do not need an account on the remote machine. Files created on the remote machine are owned by the user running gcp if the user has an account. Otherwise, the files are owned by the daemon. To copy a file, you must have read permission on your local system. You cannot use the gcp command to copy from one remote host to another; either the source or the target files must be local.

By default, you can copy a file only into the /tmp directory. An error occurs if you attempt to copy a file into any other directory. This is an intentional restriction to protect you from inadvertently filling other directories or overwriting needed files.

If you regularly copy files into other directories, the restriction becomes inconvenient. To work around the restriction, edit the last line of the /fnsw/local/lib/glogin/.security file. For example, to add write permission to the /fnsw/local/fs/FileNet/forms directory, the last line of the .security file would look like this:

0 w /tmp,/fnsw/local/fs/FileNet/forms

Note

The gcp command is not available for Windows Server Image Services systems.

The following table describes each gcp option:

Option	Action
-i	Interactive. Use this option to display a prompt with the name of the file whenever the copy will overwrite an existing file. You can respond y (yes) to overwrite the file; any other response prevents overwriting.
V	Verbose. Use this option to continuously display the number of bytes transferred so far.
n#	Sets destination network number for copying to a remote host on another network. Uses 512-byte packets.
-1	Specifies large, 1024-byte packets (the default)
-s	Specifies small, 512-byte packets

ixdb stat

The ixdb_stat command, which you run on the index server when the FileNet system is not in use, gathers statistical information about the index database and WorkFlo queues. To examine how the system is using space in the index database, see "spacerpt" on page 525.

The syntax of the ixdb_stat command is:

Enter the command name with no parameters for a help text display.

The following table describes each ixdb_stat option.

Option	Action	
-d	Gathers only document statistics	
_f	Gathers only folder statistics	
-a	A shorthand notation for specifying both -d and -f	
-v	Continuously prints the number of database rows examined so far (after each 10,000 rows)	
<output file=""></output>	User-specified name of the output file, which defaults to:	
	/fnsw/local/tmp/ixdbs_ <mondd> (UNIX)</mondd>	
	\FNSW_LOC\tmp\ixdbs_ <mondd> (Windows Server)</mondd>	
	where MonDD is the current month and day.	

Use the less or more commands to view the report; you can also print the report. A partial sample report using the -a option and displayed with the less command follows.

```
moorea(fnsw)/fnsw/local/tmp> less ixdbs Jan28
                   INDEX DATABASE STATISTICS
Date of Report : Thu Jan 28 12:17:13 2003
DOCUMENT STATISTICS :
Total Document: 43933 lowest doc id: 100004 highest doc id: 20015894
Documents with Archive Dates: 16 Archivable Today *: 16
Documents with Delete Dates: 26 Deletable Today *: 0
Closed Documents: 16
   * unless document is filed in a folder
FOLDER STATISTICS :
Total Folders: 0 Total Document Filing: 0
Folders with Archive Dates: 0 Archivable Today: 0
Folders with Delete Dates: 0 Deletable Today: 0
Closed Folders : 0
DATA DICTIONARY STATISTICS :
Total Document classes: 43
Total User Indexes: 61 inverted: 13
Document Class Total Documents Index
                               0 index2
old dcl1
                                  index1
                               0 index1
apr1601
```

		index2
apr1602	0	index1
		index2
apr1603	0	index1
		index2
apr1604	0	index1
		index2
apr1605	0	index1
		index2
apr1606	0	index1
		index2
apr1607	0	index2
		index1
apr1608	0	index1
		index2
apr1609	0	index2
		index1
apr1701	0	index1
		index2
apr1702	0	index1
	_	index2
apr1703	0	index1
1504	0	index2
apr1704	0	index1
0.1	•	index2
apr21	0	
test1	16	test
		CostIndex
	0	Menu
test3	0	
dfd	0	
test5	0	
apr2201	0	
apr2601	0	index1
2611	0	index2
apr2611 dsfj	0	
usij	0	

dsdf	0	
Akte	2	
MyTest	0	SSN
new	7	
new1	0	
DVT_docclass1	22482	
DVT_docclass2	20320	
DVT_docclass3	29	
HP_Erase_4X	22	
HP_13G1K_class	6	
HP_4X	0	
HP_4X_test	5	
ANSI_4X_Erase	5	
Remote_8X	986	
OSARLESS_REMOTE_8X	8	
OSARLESS_REMOTE_4X	3	
1kw52	17	
Erase_2X	1	
philipclass6	23	
HP4_LK_class	1	

Index	Non-Null Values	Number of Document Classes
index2	0	15
index1	0	15
apr041401	0	0
apr1402	0	0
apr1403	0	0
apr1404	0	0
apr1405	0	0
apr1406	0	0
apr1407	0	0
apr1501	0	0
apr1502	0	0
apr1503	0	0
apr1504	0	0

apr1505	0	0
apr1506	0	0
apr1507	0	0
indexapr1607	0	0
indexapr1609	0	0
indexapr1701	0	0
apr1702	0	0
apr2001	0	0
apr2202	0	0
test	0	1
test1	0	0
problem	0	0
pr	0	0
prob	0	0
apr2302	0	0
indexapr2602	0	0
indexapr2603	0	0
indexapr2604	0	0
indexapr2605	0	0
indexapr2606	0	0
indexapr2607	0	0
apr2610	0	0
indexapr2610	0	0
indexapr2710	0	0
apr2711	0	0
apr2710	0	0
indexapr2720	0	0
indexapr2730	0	0
May11index2	0	0
May1101index2	0	0
apr2002	0	0
Aktenregister	0	0
StatusDokument	0	0
WorkingSetKennz	0	0
Scandatum	0	0
Sonderzuordnung	0	0
Loeschdatum	0	0

Bestandsvorgang	0	0
Stammaktenzeichen	0	0
Aktenzeichen	0	0
SSN	0	1
test11	0	0
rtr	0	0
date	0	0
Menu	0	1
CostIndex	0	1
StringIndex	0	0

WORKFLO QUEUES STATISTICS:

Workflo Queue Name Number of Rows

Total # of Workflo Queues: 0

NOTE: Removing a document class that has a non-zero number of documents associated with it will cause those documents to be unretrievable. Removing a user index during the conversion process will cause any documents with non-null values for that index to lose those index values forever.

Done Thu Jan 28 12:17:21 2003

less

The **less** command (UNIX systems only) reads a file forward and backward.



less <filename>

Since the less command does not need to read the entire input file before starting, less can read a large file faster than a text editor can. Enter less at the command line without a filename to see a comprehensive list of options and their functions. While less is running, you can change most of the options by using the - command.

You can use a dollar sign (\$) to signal the end of an option string. This is important only for options like **-P** which take a following string. Options are also taken from the environment variable LESS. For example, for prompting in the style of the more command you can put the command **less -m** in your .cshrc or .profile file.

seteny LESS -m

The environment variable is parsed before the command line options, so command line options override the LESS environment variable.

PRI tool

Use **PRI_tool** to manage print services and the print queue, to enable and disable printers, and to reassign print requests from one printer to another. Before you start PRI_tool, you must log on to FileNet security. You must run PRI tool on the server where print services is running. UNIX users can glogin to that server.

Type a ? (question mark) at the PRI tool prompt to get a list of options:

```
costa5(kehr)/home/kehr> PRI tool
Type '?' for help
<PRI tool>?
Type 'help <command>' for detailed help on one command.
Type 'help *' for detailed help on all commands.
Type '<command> ?' for interactive command parameter
input.
Application commands:
 cachestatus cancel
                               checkcache
 clearrequests hardcopy
                               help
 modify
           printerstatus quit
 requeststatus resumeprinter
                               systemstatus
  termoff
                termon
```

You can abort any command using Control+c. Enter **q** to exit PRI tool.

cachestatus

The cachestatus command displays the number of objects fetched for each cache and for each printer. Use this command to determine if the number of pages fetched is close to the cache full threshold, which can prevent a printer from printing.

```
<PRI tool>cachestatus
Summary statistics for caches (all printers):
  Cache 0: sys_print_cache1:costa5:FileNet
            min/max pgs=409/819, #sects=81920, pgs fetched=2
Statistics on each cache for each printer:
  Printer 0: Xerox:costa5:FileNet
      Cache 0: sys print cache1:costa5:FileNet
         min/max full=50/80, min/max pgs=409/655, pgs fetched=0
Printer 11: fax4:costa5:FileNet
      Cache 0: sys print cache1:costa5:FileNet
         min/max_full=20/50, min/max_pgs=163/409, pgs_fetched=2
```

The first section shows summary statistics for each cache, including these items:

Item	Statistic
cache	Shows the ID of the displayed cache
minimum pages	Shows the cache threshold at which the minimum pages per printer limit changes
maximum pages	Shows the number of pages the cache will hold
# sectors	Shows the number of sectors in the cache
pages fetched	Shows the number of pages for which a fetch has been initiated for the listed cache, regardless of which printer queue contains the corresponding request. This counter increments for each fetched page and decrements when a page is deleted. The counter increments before the fetch completes and, therefore, does not reflect the number of pages in the cache.

The second section contains the number of pages fetched into each cache for each printer.

Item	Statistic
cache	Shows the ID of the displayed cache
printer	Shows the displayed printer. Information displays for each printer and cache combination.
minimum fullness	Shows the minimum percent cache fullness threshold
maximum fullness	Shows the maximum percent cache fullness threshold
minimum pages	Shows the minimum cache fullness threshold as a number of pages
maximum pages	Shows the maximum cache fullness threshold as a number of pages
pages fetched	Shows the number of pages for which a fetch has been initiated without a delete for the listed cache and printer

The cache fullness affects fetches in three ways:

- If the number of pages fetched is less than minimum pages for a given cache (first section), the system fetches up to the maximum page limit specified for the cache and printer (second section).
- If the number of pages fetched is greater than minimum pages and less than maximum pages for the cache (first section), then the system fetches up to the minimum page limit for the cache and printer (second section).
- If the number of pages fetched is greater than or equal to maximum pages for the cache (first section), no additional fetching is done for any printer.

cancel

The **cancel** command cancels the indicated requests. If you include multiple options, the command cancels only the requests that satisfy all options.

Option	Description
request ID	Specifies the request ID (job number) to cancel
filename	Specifies the name of the file that contains a list of request IDs, one request ID per line
username	Cancels only requests from this user
printer name	Cancels only requests for this printer
priority	Cancels only requests with this priority

checkcache

The **checkcache** command checks for objects that don't belong in the print cache and suspends all activity by print services while checkcache runs. Runtime can range from a few seconds up to 15 minutes for systems with more than 100,000 pages waiting to be printed. If necessary, you can abort this command (Control+c).

clearrequests

The **clearrequests** command deletes all print requests and removes all objects from the print caches.

You might use this command to clear requests that have frozen a print server. For example, if the print server violates a protocol, it dumps the trace buffer to a file and logs an appropriate message to sys_log.

hardcopy

Use the **hardcopy** command to send output from your display device to a file on magnetic disk.

hardcopy [<filename>]

If you do not use a filename, this command acts as a toggle to turn hardcopy off and on. See also "termoff" on page 523.

help

Displays help text for all commands or for just the specified command.

```
help [<command>]
```

To display the list of commands, enter either a question mark (?) or help with no commands.

To display the complete help text file, enter:

help *

The help * command is most useful if you are working in a window environment where you can scroll back to see the start of the file. Otherwise, use the hardcopy command or control the screen display with Control+s and Control+q.

To get help for a specific command, enter the command with no parameters. Here is an example, using the cancel command.

```
<PRI tool>cancel
request id (CR for other options)?
name of file with request ids (CR if none)?
only requests from user (CR=any)?
only requests from printer (CR=any)?elsinore
only requests with priority (CR=any)?
```

modify

The **modify** command works with either a request ID or a file containing a list of request IDs (one request ID per line) to modify a print request.

If you specify more than one of the options that start with **from**, a request must satisfy all of these options to be modified. If a request is already printing, you can only modify printer, paper size, or overlay, or suspend the request by setting the priority to 0.

For changes after printing has started, the system cancels and resubmits the request, which can change the printing order relative to other requests.

```
modify [<request id>] [fromfile=<filename>] [fromuser=<username>]
    [fromprinter=<printer>] [frompriority=<priority>]
    [frompapersize=<papersize>]
    [paper_size={letter,legal,b,c,d,e,a0,a1,a2,a3,a4,a5,b4,b5,18x24,top,
              bottom, third, default, half letter, best
avail, 10x14, executive ]
    [scaling={normal,clipboth,exact,approx,original,center,enhanced
exact 11
    [orientation={default,landscape,portrait,no rotate}]
    [form name=<form name>] [note=<note>] [priority={0..9}]
    [printer=<object:domain:organization>] [copies=<#copies>]
    [overlay={none, firstpage, allpages}] [eject tray=<eject tray#>]
    [print time=<yyyy/mm/dd hh:mm:ss>]
    [security=<read name> <write name> <exe name>]
    [staple] [two sided] [collate] [annotations] [request header]
    [doc headers] [phone num=<phonenumstr>] [headline=<headlinestr>]
    [fax mode={coarse,fine}] [page footnote] [time footnote]
```

The following table lists and describes the action of each option you can use with the modify command.

Option	Description
<request_id>:</request_id>	Specifies the ID of the request to modify
fromfile	Specifies the name of a file containing request IDs
fromuser	Modifies only requests from this user
fromprinter	Modifies only requests for this printer
frompriority	Modifies only requests with this priority
frompapersize	Modifies only requests with this paper size
paper_size	Specifies paper size to use
scaling	Specifies scaling option to use
orientation	Specifies which edge is the top
note	Specifies an ASCII string you provide to print on each header page
priority	Specifies priority of job, 0=lowest and suspends printing; 9=highest
printer	Specifies NCH name of printer
copies	Specifies number of copies to print
overlay	Specifies overlay no pages, the first page, or all pages. When specifying first page or all pages, the first page contains the overlay data
eject_tray	Specifies the number of the output tray where pages are to be delivered
print_time	Specifies time to start printing using the system's default date format. For example, if your system is configured with a date mask of "day mon dd hh:mm:ss yyyy," you would enter the date using a format as shown in this sample: Mon Nov 15 12:30:00 1998).
security	Assigns security names for read, write, and append/execute
collate	Collates pages. Prints entire copy of each document before beginning second copy.
annotations	Prints annotations
request_header	Uses a header page for the print request
doc_headers	Uses document header pages
phone_num	For fax requests: specifies phone number to dial

Option	Description
headline	For fax requests: specifies text of headline to print at top of each page
fax_mode	For fax requests: chooses coarse or fine for resolution

printerstatus

Use the printerstatus command to get a long or short report of the status of one or all configured printers and fax servers.

printerstatus [<printer_name> [nonfax] [faxonly] [long]]

Entering the command with no parameters displays a short report on all printers and fax servers:

<pri_tool>printerstatus # printer print name request</pri_tool>	fetch request	requests queued	pages queued	idle time
0 RIC20_4:costa5:FileNet		0	0	9632
1 Xerox:costa5:FileNet		1	1	suspended
2 Versatec:costa5:FileNet		0	0	9632
3 fax1:costa5:FileNet 0		0	0	down
4 fax2:costa5:FileNet 7467 50 50 down		down		
5 fax3:costa5:FileNet	7528	2	2	down

Following is a long status display for Xerox:

<PRI tool>printerstatus Xerox

Printer 1: Xerox:costa5:FileNet

Status: suspended

Print error: (n.a. -- suspended)

Request printing: none

Embedded migrate: (n.a. -- suspended)

Next migrate: Time since print: 3567 Requests queued: 1 Pages queued: 1 Pages printing:

Avail papersizes: (n.a. -- suspended)

Config papersizes: letter, legal, a4, b5, top, best avail, default

The following table describes the information in the fields of the printer status display.

Field	Content
Status	Shows the state of the printer: unknown, down, suspended, needs attention, needs service, available, or redirected. If the status is redirected, the destination printer name also displays.
Print error	Shows the error tuple for error that is preventing printing or 0 if there is no error. Use fn_msg to see the text for the error.
Request printing	Shows the ID of the request that is currently printing, the number of pages printed for this request, and the total number of pages in the request
Embedded migrate	Shows the ID of the document being migrated due to an embedded document reference and the name of the document service receiving the request
Next migrate	Shows the name of the service and the document ID if the printer is waiting for a page. Otherwise, this field shows done .
Time since print	Shows the time elapsed in seconds since the last page was printed
Requests queued	Shows the number of requests queued to this printer
Pages queued	Shows the number of pages queued to this printer
Pages printing	Shows the number of pages submitted to the print server but not completed
Avail papersizes	Shows the currently loaded paper sizes
Config papersizes	Shows the complete list of paper sizes configured for the printer

requeststatus

Displays the status of one or more print requests.

requeststatus [<request_ID>] [user=<username>] [long] [docs] [priority=<priority>] [printer=<printer name>

Each option you can use with the requeststatus command is listed in the following table, with a description of the result of each option when used with the command.

Option	Description
request ID	Specifies the ID of a print request. If you do not specify an ID, the display includes all requests that satisfy other options.
username	Specifies the logon name of a user whose print requests you want to view
long	Displays detailed information. If not specified, the display is only summary information.
docs	Generates a list of the items being printed: document ID for documents and object ID/system serial number for cache objects
priority	Displays requests for the specified priority:
	0=lowest and suspends printing
	9=highest
printer name	Displays requests queued to the specified printer. Requests appear in the order they will be printed. (When you do not specify a printer, requests display in ascending request ID order.)

The following is an example of a printout when a request ID is not specified.

<pri_too< th=""><th>l>reque</th><th>ststatus</th><th></th><th></th></pri_too<>	l>reque	ststatus		
request	total	pages	status	printer
id	pages	printed		name
7466	1	0	queued	Xerox
7467	1	0	printing	fax4
7468	1	0	fetching	fax4
7469	1	0	queued	fax4

The following example printout is a long report for a single request ID.

```
<PRI_tool>requeststatus 5003 long
request_id=5003, request_status=queued, print_err=00000000, fax_request=f,
submit_time=93/12/8 16:00:58, print_time=n.a., done_time=n.a., copies=1,
priority=4, paper_size=dont_care, collate=f, two_sided=f, annotate=f,
req_header=t, doc_headers=f, scaling=normal, orientation=default,
printer=Xerox:costa3:FileNet, user=SysAdmin:costa3:FileNet, form name='',
note='', eject_tray=0, total_pages=1, pages_printed=0
```

The following example printout shows the information given with the docs option.

```
<PRI tool>requeststatus docs
reg id=5003, total pages=1, status=queued, printer=Xerox
doc id=30093216, 1stpg=1, lastpg=1, service=DocServer:costa3:FileNet
req_id=5004, total_pages=1, status=queued, printer=Xerox
doc id=30093216, 1stpg=1, lastpg=1, service=DocServer:costa3:FileNet
```

resumeprinter

The **resumeprinter** command starts a printer that was previously suspended or redirected.

resumeprinter <printer name>

systemstatus

The **systemstatus** command gives overall printer subsystem status information. For example:

```
<PRI_tool>systemstatus
Total requests queued: 53 Total requests printed: 0
Total pages queued: 53 Total pages printed: 0
```

termoff

The termoff command turns off output to the terminal. Use this in conjunction with the hardcopy command to get large listings into a file without waiting for terminal I/O (see "hardcopy" on page 514).

termoff

termon

The termon command turns on output to the terminal after using the termoff command.

termon

spacerpt

The **spacerpt** program reports on the space use of the FileNet tables and associated B-trees in an Oracle, DB2 or Microsoft SQL Server RDBMS database. By default, spacerpt reports on all FileNet tables, including FolderView tables you store on the index server. Run spacerpt to monitor the space in the index database and to monitor the space in a WorkFlo Queue database. A table name beginning with the prefix WQM is a WorkFlo queue.

Note

To run spacerpt on a Remote Oracle server, see "Running Spacerpt on Servers with Remote Oracle Databases" on page 529.

To establish a baseline for space use, run spacerpt weekly. Then, depending on the rate of growth observed, run it more or less frequently to monitor the database and plan for magnetic disk expansion.

Limit the report by naming specific tables as arguments:

```
spacerpt [tablename ...]
```

To get a detailed report on all extents (Oracle, DB2) or indexes (MS SQL Server):

```
spacerpt -x [tablename ...]
```

You can also limit the report to specific owners of the tables:

```
spacerpt –u f sqi (to see just the FolderView tables)
```

spacerpt –u f_sw (to see all tables except FolderView)

The program outputs to stdout (the console screen).

If you are not working in a scrollable window, you may want to redirect the output to a file that you can view with the less utility.

spacerpt > /tmp/db.out

The following abbreviated sample illustrates results for a DOCTABA created using the Oracle RDBMS. If DOCTABA was created with MS SQL Server, some headings would differ and names would be in lowercase.

corona(fnsw)/ho	01/2	9/03 SPAC	E RE	PORT		C	orona		
SYSTEM		2047 4095			dev/1/ dev/1/				
sum		6143			~+ £	_			
Tablespace		Free space (K		_					
SYSTEM ROLLBACK SEGMEN		_	62		1975	0			
Rollback Segmen	t						_	Last ext (KB)	
RS0				2	9999	0	100	100	200
RS1				2	9999	0	100	100	200
RS2					9999		100		
RS3					9999	0	100		200
SYSTEM				4	9999	0	50	50 	200

sum							1000
TABLE SPACE							1000
Table (KB)					_	Last ext (KB)	Total Space
 DOCTABA INDEX SPACE		1	121	0	10000	10000	10000
- 11	- 1				_	Last	
Table (KB)	Index	use	d exts	incr	ext (KB)	ext (KB)	Space
DOCTABA	DA_A32	1	9999	0	5000	5000	5000
	DA_ARCHIVEDATE	1	9999	0	50	50	50
	DA_DELETEDATE	1	9999	0	50	50	50
	DA_DOCNUMBER	1	9999	0	5000	5000	5000
	DA_PURGEDATE	1	9999	0	50	50	50
sum							10150

The following table describes the columns and fields in the spacerpt report.

Field	Description
Tablespace	Shows table space name
File #	Shows the sequential number of each file within the database
Size	Shows the size of the file in kilobytes
Name	Shows the full path name of each index database data file
sum	Shows the sum of all the space of all files listed
Free Space	Shows the amount of space not allocated to any table, rollback segment, or index
Largest Free Extent	A database table is divided into regions called extents . Extents contain the actual indexing information that users enter. This field reports the size in kilobytes of the largest free (unused) extent for the table space.
Rollback Segment	Shows the name of the rollback segment
Table	Shows the table name
Exts used	Shows the total number of extents used by the table so far
Max exts	Shows the maximum number of extents that can be allocated for the table
Pct incr	Shows the percentage by which the size of a newly allocated extent is increased over its predecessor
Largest ext	Shows the size in kilobytes of the largest extent allocated for this table
Last ext	Shows the size in kilobytes of the last extent allocated for this table
Total space	Shows the amount of space, in kilobytes, allocated to the table

The part of the report entitled Index Space provides information about the index space for each table. All field descriptions are the same as given above. One additional column, Index, shows the internal name of each index.

Running Spacerpt on Servers with Remote Oracle Databases

If the Oracle databases are located on a remote Oracle server, you cannot run spacerpt directly because Oracle OS authentication prevents it. Instead, you need to modify two script files and login to sqlplus to get space information.

1 Use your preferred text editor, such as **vi**, to modify these two files:

> /fnsw/oracle/spacerpt summary.sql /fnsw/oracle/spacerpt extended.sql

The first line of each of these two files is:

/ as sysdba

2 Remove this line from each file. Exit and save your changes.

Now you can run the **spacerpt** scripts successfully on the Image Services system with remote Oracle databases.

3 Login to sqlplus to run the scripts:

sqlplus

- 4 When you're prompted, enter the user name f_maint and f_maint password.
- 5 To run **spacerpt**, enter the following command at the SQL> prompt:

SQL> @/fnsw/oracle/spacerpt summary.sql

ssn

This Image Services command displays the system serial number.

ssn

stdocimp

Use the **stdocimp** utility to import just the documents that did not commit successfully using the Import Documents from Media function (see "Import Documents from Media" on page 384). Before reimporting, you must determine the cause of the failed committals and correct the problem.

Typing **stdocimp** with no parameters displays help text for the utility. The stdocimp program has the following syntax:

```
stdocimp [-nonexactclass] [-noinsertdoctaba]
   [-security {none | doc}] [-updatesnt] [-noworkfloqueue]
       [-nodeleteddocs] [-bothsides] [-highpriority]
           [-optdiskfile <optdiskfile>] [-skipcount <skipcount>]
              [-docidfile <docidfile>] [-redo <redojobnum>]
                  [-ssn <ssn>] [ignoredeleteupdate]
                      [overwritedoctaba] <surface ID>
```

The only required entry is the **surface_ID**, except when using **-redo**.

Note

If stdocimp encounters an error (such as mismatched document class) while importing a document into DOCTABA, stdocimp corrects an inconsistency in the databases based on DOCTABA in the following manner:

If the document exists in DOCTABA but not in DOCS prior to the import. the document is imported to the DOCS table.

If the document exists in DOCS but not in DOCTABA prior to the import, stdocimp deletes the document from the DOCS table.

The following table describes the stdocimp options.

stdocimp options

Option	Description
-nonexactclass	Imports documents when the indexes in the document classes do not match exactly. If an index in the imported document is not defined in the index database, the index is dropped. If an index defined for the document class is missing from the imported document, the index is set to null. An index not matching the type in the database is dropped.
-noinsertdoctaba	This option is mutually exclusive with the -overwritedoctaba option.
	Inserts a record into the permanent database, but does not insert the record into the index database. Use this option only on systems that do not use FileNet's index database to store indexes.
-security {none doc}	Sets security to the value from the document on storage media if you specify –security doc. Specifying none sets the security to (ANY-ONE). The default is to use the security defined for the document class (doc).
-updatesnt	Updates the scalar_numbers table if the document ID to be imported is greater than the next available document ID. This option allows you to reinsert documents into a database if the database on magnetic disk has been back dated by a restore.
-noworkfloqueue	Does not insert the document into a WorkFlo queue if a queue is defined for the document class.
-nodeleteddocs	Imports documents that are partially imported—with a record in either the permanent or index database, but not both. This option allows you to synchronize the two databases.
-bothsides	Imports documents from both sides of storage media. The option is ignored if you specify a document headers file (–optdiskfile).
-highpriority	Runs at high priority (same as retrievals).
-optdiskfile <optdiskfile></optdiskfile>	Specifies the name of the document header file listing the documents to be imported. If you aren't sure of the name, you can display a list of document header files using the Background Job Control program.

stdocimp options, Continued

Option	Description
-skipcount <skipcount></skipcount>	Specifies how many documents to skip before starting to import them. For example, if skipcount is 100, the first 100 documents are not imported unless they are explicitly specified in a document ID file list.
-docidfile <docidfile></docidfile>	Specifies a magnetic disk file name where each line is [ssn] doc_id. If an SSN (system serial number) is present, the SSN and document ID must be separated by spaces. If an SSN is not present, the program imports the specified document regardless of the SSN. If document IDs are reassigned during import, the document IDs in this file are the IDs on storage media, not the IDs on the importing system.
	If you specify both a skipcount and docidfile, a document that satisfies either condition is imported. For example, if skipcount is 100, and several documents in the first 100 are imported by docidfile, the skipcount option still begins processing with the 101st document.
	The –docidfile option requires 10 bytes of memory per document in the docidfile.
	Any documents not imported because they were not present on the media are logged with an error message indicating they were not found.
-redo <redojobnum></redojobnum>	Imports documents not imported by the background job number redojobnum. You cannot use this option to redo an import job that used a file list of document IDs. Using –redo sets the options that were previously set for that job. You can override any of these options by entering additional options later on the command line.
-ssn <ssn></ssn>	Specifies the system serial number of the surface containing documents to be imported if the SSN is for an incompatible system.
-ignoredeleteupdate	Allows short descriptors that have been marked Deleted by the stsur- fupdate tool to be imported to a system's database. Normally a short descriptor that's marked on an optical disk cannot be imported into a system's index database.

stdocimp options, Continued

Option	Description
-overwritedoctaba	Allows stdocimp to replace index information in doctaba with matching index information from the short descriptors on the optical media.
	However, if a value exists in the index database and the corresponding value for the same document on the optical disk is null, the Index dataset value will NOT be overwritten.
	This object is mutually exclusive with the -noinsertdoctaba option.
surface_ID	Specifies the surface ID containing documents to be imported. Do not use this option if you use the –redo option.

Printing

This chapter describes the printing process in general terms. You use a PC running your desktop application (not the server console) to print image documents or modify the print queue. Use a PC workstation to make any changes to the print job through the print queue. See "PRI tool" on page 510 and the "Printing" chapter in your desktop application's user's guide.

For details, see the following topics:

- "How Printing Works" on page 535
- "Security" on page 546
- "Troubleshooting" on page 547

How Printing Works

The printing function on an Image Services system consists of three major sections: printer applications, print services, and printer imaging.

At the highest level are the printer applications that contain the user interface, terminal display, and operator input.

The middle level of the print software is print services. This layer is the body of software responsible for determining the order to process print requests, migrating documents from storage media to magnetic disk cache, and modifying and canceling print requests. Print services also maintains status on print requests and printers, making this information available to the print application layer, which can display current printer activity.

Print services does not know about the format of any page or annotation submitted for printing. The task of print services is to get the data into the cache in the user-specified format.

The lowest level of the print software is printer imaging. This software accepts a cache object as input data, reads this data into memory, and transforms it into a form the printer can use. It then issues an I/O request to the printer.

You can print margin notes and highlights directly on the image page; annotations appear on a second page. When printing highlights, the printer uses shading. It does not attempt to simulate the colors associated with either a margin note or highlight.

Print Services Database

The print services database consists of four MKF tables in the transient database: print requests, print options, print docs, and print SVCS.

The required data for a print request is in the print requests table, along with the status of a request, statistical information about a request, and information about the document or cache object to be printed (if only one).

The print options table contains the user-specified print options. If the operator does not specify any options, the table contains only a print request ID.

The print docs table contains information about which documents or cache objects to print for this request (if an operator specifies more than one document or cache object).

The print_svcs table maps document and cache service requests to a two-byte ID number.

Printer Selection

If you do not choose a printer or fax server, print services chooses from the list of devices it knows about based on the following criteria:

- If you requested a printer, print services eliminates all fax servers. It eliminates print servers when you request a fax server.
- It eliminates any available print device that is not operational.

It eliminates any print device not capable of printing as requested (for example, on the requested paper size or from the requested paper tray).

From the list of remaining print devices, print services selects the device with the least number of pages queued to print. If all the devices are the same, which is the case when none of them is busy, it assigns the job the first listed print device.

Note

Since the system sees each UNIX file as one page, it reads a printer processing a 100-page UNIX file as printing one page.

You can reassign individual waiting print requests to another printer. See the "Printing" chapter in your desktop application's user's guide.

Use the PRI tool redirectprinter subcommand to move all gueued jobs from one printer to another. Call your service representative for the password you need to use this command.

Request Ordering

Print order is partially determined by the sequence number in the print options table. The sequence number increments for each new request. Assigning a new priority or changing printers generates a new sequence number, causing the print request to go to the end of the queue for that priority and printer. No other modification causes a change in the print order.

Variable-length Pages

You can have your system configured to treat each page of a document separately when selecting a paper tray.

The Best Available print option enables the printer to choose the most appropriate paper size for each page, while keeping all the pages of a document together. The system treats each page separately and determines the best available paper size. When not configured for this option, the system looks at the first page of the document to determine the best available size, then uses that size for the entire document.

Ask your service representative to edit the appropriate print configuration file to use this feature.

Print Services Initialization

When print services initializes, it sends a message to each print server. Receiving this message re-initializes each print server, while allowing print services to keep a table of the available print devices. If it does not receive a response from a print server, print services marks its status as **unknown**. When a print server initializes, it sends a message to print services indicating that this server has rebooted. When print services receives this message, it reissues all outstanding page print requests to the server and marks the status of the print server as **up** to begin page print requests to the server.

The configuration information describing the printers assigned to print services resides on the print services station (the server on a combined server system). Having this information available to print services at all times allows the system to gueue requests to a print server even if the print server is down, preventing the need to reject requests due to no known compatible printer.

Note

To change the print service a particular printer is associated with, simply delete the printer from the old service, then add it to the new service. Deleting and adding printers is done through the Procedures tab of the System Configuration Editor (see the on-line help for System Configuration Tools).

Application and System Print Caches

A print service uses two types of cache: application print cache and system print cache. Each print service has one application print cache, used by applications that print data other than image documents. Print services uses the system print cache to retrieve documents from storage media. The system transfers the document only once, which helps if the server has print services but is not a storage library server.

If more than one FileNet system uses one printer, setting up a system print cache, print service, and printer all on the same server is the most efficient way to use the printer. However, such a server requires a magnetic disk.

On a system with more than one print service, only one print service can contain a system print cache on each storage library server. All other print services must have one system print cache, preferably on the server with print services. Only one print service can reside on each server.

Caching Strategy

The system retrieves documents directly from the storage media to the print cache. While there is no prefetching to the page cache, the system checks page cache first and copies the document from page cache to print cache, if possible.

The system uses the print request ID to identify objects in print cache instead of using reference counts, system serial numbers, and document IDs. Using the print request ID eliminates the need to inspect all objects in the cache if you need to restart print services.

When printing documents from the local system, print services uses a default cache name that is mapped to the real cache name. Thus, with multiple storage library servers, your service representative can set up a default print cache on each storage library server and the data transfer process can always write the image data into a local cache. Therefore, it transfers the image data only once across the network during printing, regardless of system configuration.

Configuring print services on a remote server enables it to print objects from local page cache, rather than retrieving them from across the network.

The caching strategy uses a cache-full recovery algorithm for one-step retrieval. When a cache-full error occurs, print services must delete from cache any objects that could potentially prevent the First In First Out (FIFO) order processing. Print services also reduces the number of retrieval requests it has outstanding so that the condition cannot occur again.

When deleting objects after a cache-full error, print services first tries to delete any cache objects that correspond to printed pages for jobs that specified multiple copies. It deletes cache objects after printing one copy of the page, but does not normally delete cache objects until it has processed the whole request for multiple-copy jobs. After deleting all multiple-copy, printed pages, print services inspects each printer and deletes a percentage of the lowest priority requests for each one.

Printing UNIX Files



Use XPR print at the command line to print UNIX files. To print reports, such as those available in Database Maintenance and SLC, select the print option from within the application. To print UNIX files, the window manager must be running. If you did not answer yes to the prompt to start X-windows, you can start the window manager with either of these commands:

xinit


mwm &

You must be logged on as either root or fnsw, or you must log on to FileNet security using the fnlogon program. See "UNIX Users Logon" on page 469 for more information about fnlogon. Type the command with or without a filename parameter.

XPR print [<filename>]

Print File or Report Dialog Box

Using the XPR_print command displays the same window you see when printing from an application, as shown in the following example.



The following table shows the options in the Print window.

Option	Use
Data	If you did not enter a file name with the XPR_print command, this area is blank. Click in the Data box and enter the name of the file you want to print.
Service	To select another configured print service, click the down arrow and select from the list.
Device Type	Printer is automatically selected. Click FAX if appropriate.
	Click the Additional Print Servers button to send your print request to a printer on another system or to a different print server on your own system. When you click this button, the list of print servers in the Service field includes all additional print servers beyond the default. Select the appropriate print server, then select the appropriate printer in the Printer field.
Printer	Click the list box to select an alternate printer for the default print server. To review the status of all printers and fax machines configured on this server, click the Status button. The list shows the status (Queued, Printing, etc.) and the number of print requests for each printer. To see additional printers on other print servers (local and remote), check the Additional Print Servers box.
	Changing the device type does not affect the name in this box. You must choose an appropriate printer for the device type.
Priority	If the default of 4 is not appropriate, click a higher number to print sooner or a lower number to print after higher priority jobs. If you click 0, the job is suspended until you select a higher number.
Paper Size	To choose a different value, click the down arrow and choose another item on the list. Be sure to choose an appropriate size for the selected printer.
Copies	If you need more than one copy, change the 1 in the text box to the number you need.
Note	Enter up to 40 characters that you want to print on the job header page.
Options	Turn the job header (cover page) on and off by clicking the button.
Printing the Report	Click OK or Cancel after filling out the items in the window. After submitting the request, the Data field clears and you can enter another file name or click Cancel to close the window.
Modifying a Print Job	You will need to use a PC workstation to make any changes to the print job through the print queue. See the "Printing" chapter in your desktop application's user's guide. See also "PRI_tool" on page 510.

Interrupting a Print Job

When printing a large document, the software marks a checkpoint every ten pages. If the job is interrupted (for example, in the case of a server failure), the system resumes printing at the last checkpoint.

A failed print job is marked with a **Failed** status. The system does not change its priority; neither does it automatically restart jobs that fail.

While a job is printing, you can change its priority to 0. This suspends the job so that you can modify it as desired.

You can cancel any print job from a PC workstation. You do not need to stop the job or wait for paper; you can cancel actively printing jobs. A canceled job stays in the queue for ten minutes so you can easily restart the job if you change your mind.

Security

While printing, print services checks several security items:

- Logon name, terminal, and password of the individual logged on
- Groups assigned to the printer (considered a terminal)
- Access restrictions on documents
- Access restrictions on annotations
- Access restrictions on the print request
- Access restrictions on other objects in cache

Print services checks the access restrictions on documents, annotations, or other cache objects to insure that both the user requesting the print and the printer itself has access to the object. Ensuring the printer has access to the document prevents printing sensitive material at a non-secure site.

Access restrictions on the print request determine who can modify or cancel the print request. See Chapter 3, "Security Administration," on page 194 for information about setting up security.

If a printer is not printing and the problem is not obvious, follow the steps below to determine the cause.

- 1 Log on to FileNet security and run PRI tool.
- 2 Enter **p** to display the short form of the printer status report.

The Requests Queued and Pages Queued columns should be nonzero for the printers you expect to be busy.

<p< th=""><th colspan="6"><pri_tool>p</pri_tool></th></p<>	<pri_tool>p</pri_tool>					
#	printe	rprint	fetch	requests	pages	idle
	name	request	request	queued	queued	time
0	0 RIC20_4:costa5:FileNet 5054			2	2	down
1	1 Xerox:costa5:FileNet			0	0	suspended
2	2 Ricohprnt:costa5:FileNet 0			0	1281	51

The Idle Time column shows the number of seconds this printer has been idle, or the words suspended, redirected, down or unknown. If the idle time is less than the time required to print a page, or if no requests are queued, the printer is working normally. Use the resumeprinter command to restart a suspended or redirected printer. Otherwise, continue to the next step.

3 Enter **p** printername> long to display the long form of the printer status report.

```
<PRI_tool>p RIC20_4 long
Printer 0 :
                 RIC20 4:costa5:FileNet
Status:
                  down
Print error: <181,0,13>
Request printing: request id=5054, total pages=1, pages printed=0
Embedded migrate: (n.a. -- down)
Next migrate:
Time since print: 128348
Requests queued:
Pages queued:
                   2.
Pages printing:
                   0
Avail papersizes: (n.a. -- down)
Config papersizes: letter, legal, b, a3, a4, a5, b4, b5, top, bottom,
third, half letter, best avail, 10x14, default, dont care
```

- If the printer is down, look at the Print Error field. The error tuple normally indicates a network failure or a failure of the print server. Check that the print server is up and running and that the network is functioning correctly.
- If the idle time of the printer is high, but the printer is not down, then look at the Pages Printing field to see how many pages were submitted. If this number is nonzero, then the printer has not reported a completion status.
 - a Check the Print Error field to see if an error is preventing the printer from continuing, and check that the Status field is available and not needs svc or needs attn.
 - b Also check the Embedded Migrate field. If a document is listed here, see why this document can't be retrieved from the storage library server.
- If the Pages Printing field is zero (nothing has been submitted to the printer), check the Next Migrate field. If a document is listed here, then check the storage library server to see why the document can't be retrieved. If no document is listed in the Next Migrate field, proceed to the next step.

Enter **s** to display the system status. 4

At the end of the display, you might see these messages:

Message	Significance
Cancel all requests in progress	The Clear command in PRI_tool is running.
Cache full recovery in progress	The estimated page size in the print_config file is wrong. The print service is deleting objects from cache and reconfiguring itself.

In either case, the printers should resume normal activity after a short while.

5 Enter cachestatus to display the summary statistics for caches (all printers).

If the pgs fetched value is close to the max_pgs value, then print services will not fetch additional pages into that cache until other pages are printed and deleted. You can configure print cache so that one printer can fill up all space in all caches, or to dedicate cache space to each individual printer. If your system does not dedicate cache space to each printer, printing may run faster. However, when a single printer receives a large number of requests, pages for that printer use up all cache space and other printers become idle.

6 Enter **r printer=<printername>** and verify that requests queued to the printer do not have a status of waiting.

Requests that are **waiting** were submitted with a delayed print time and are not supposed to be printing at the current time.

Look in the error log for messages that start with PRI. 7

The message, "PRI: Process aborted due to previous error," indicates that an error occurred and the print subsystem is hung. Report this message to your service representative. You may receive additional requests for information from your service representative.

- 8 If a malfunction occurs, try one of the following solutions (for details on the PRI tool commands, see "PRI tool" on page 510):
 - Use the **resumeprinter** command in PRI tool for a printer that is down. If successful, the printer will become available.

The system polls a down printer periodically (every 5 or 10 minutes). If the problem is fixed, the printer will come back on-line. The resumeprinter command puts the printer on-line immediately.

- With the **redirectprinter** command, you can redirect print requests to another printer. Call your service representative for the password you need in order to use this command.
- If a printer has stopped printing because it's out of a certain paper size and no more of that paper is available, you may be able to substitute another paper size.

Use the **modify** command in PRI tool to make the paper size change. You can change all requests for a given printer and paper size by entering the following command (all on one line):

modify fromprinter=<printername> frompapersize=<oldpapersize> papersize=<newpapersize>

Reboot the print server. Print services uses a reboot recovery sequence to synchronize the print server with print services.

Digital Document Transfer System (DDTS)



The FileNet Digital Document Transfer System (DDTS) is available only on a UNIX-based Image Services server. It is not available on a Windows Server system.


Overview

The FileNet Digital Document Transfer System (DDTS) consists of the export program **ddtsexp** and the import program **ddtsimp**. Using these programs, you can export and import documents in either of two formats, FileNet format or Computer-aided Acquisition and Logistic Support (CALS) format, using either tape or disk.

Use the FileNet format to transfer documents from one FileNet system to another. DDTS does not export margin notes or highlights.

The CALS format, widely used in government and industry, may be common ground for exporting documents before you import those documents into a non-FileNet system, or vice versa.

DDTS can export FileNet image documents in the CALS data raster format or in FileNet format, either of which is suitable for import into other FileNet systems. The FileNet-specific options for export and import of FileNet documents are not part of the CALS standard.



Using DDTS to export or import documents involves these steps:

- 1 Create a command file.
- 2 Select options.
- 3 Run the export or import program.

Each file you specify in the list at the end of the ddtsexp or ddtsimp command consists of a structured set of options for exporting or importing FileNet documents. Some of these options provide information required by CALS. Other options are FileNet-specific; you can use them only to export documents for later import into another FileNet system or import documents previously exported from another FileNet system.

You can use strings with either single or double quotes in DDTS command files. Each double quote (") in a double-quoted string must be preceded by a backslash (1). Likewise, each single quote (1) in a singlequoted string must be preceded by a backslash (1).

DDTS commands are free form: they ignore blank spaces. Command words are not case sensitive. A number sign (#) in the first column of a line indicates a comment line.

For details on exporting files, see "Export" on page 565. For details on importing files, see "Import" on page 574.

To abort an export or import.

- If the process is running in the foreground, use Control+c to abort.
- If the process is running in the background, use:

where <pid> is the process ID number. Check the Task Manager's PID column for the process ID number (see "Check the Active Processes" on page 330).

If the **kill –31** command fails, use the unconditional kill option:

Be sure to use the correct <pid> when using the unconditional kill option.

If you use DDTS repeatedly for similar exports and imports, you can save time and reduce errors by creating a script to streamline and automate the process. See "Creating a Script" on page 583 for more information and a sample script.

Tape Drive Support

DDTS supports 8mm, 1/4-inch (QIC), and 1/2-inch, 9-track tape drives attached to an Image Services server. If you are exporting documents in CALS format for future import to a non-FileNet system, be sure to use a tape drive and a tape density that is compatible with that of the destination system.

An export job fails if it exceeds the length of the tape. The following table shows the default tape capacities assumed by the software.

Tape Format	Capacity
8mm (assumes 112 meters)	4200 MB
1/4-inch	525 MB
1/2-inch (9-track reels)	2400 feet X density

If you use a shorter tape with less capacity, prior to running DDTS you must set the TAP MB environment variable to specify how many megabytes the tape will hold. For example, in the Korn shell you could enter:

export TAP_MB=2250

An 8mm tape that is 54 meters holds approximately 2300 MB.

An 8mm tape that is 15 meters holds approximately 600 MB.

For 9-track tape reels, multiply the length of the tape times the density. For example, 2400 feet X 6250 bpi is about 165 MB. If your tapes are 2400 feet, you do not need to set the variable since the software determines the density and sets the appropriate value. For shorter tapes, you must set new values.

Note

Image data is stored in a compressed TIFF format; DDTS does no additional compression.

Maximums

An export or import is limited to 999 documents, each of which can have up to 999 pages. If you are using tape, you may exceed the tape volume limit before you reach the document limit. If you are exporting from a file or importing to a file, you may also be limited by the available disk space for the file. The CALS format limits an import or export to three volumes; the FileNet format imposes no limits.

Performance

The tables below list the results of minimal performance testing so you can make estimates for planning purposes. Performance on your system may vary significantly.

In a series of tests, we imported and exported 400 one-page documents, where each image was approximately 100 KB in size. When each test began, the system had no other activity, the document was in cache, and the tape was in the 8mm tape drive. These are our results:

DDTS Timing Estimates

Export	Time Elapsed (min:sec)	Seconds Per Document
1st run	7:34	1.13
2nd run	7:24	1.11
3rd run	7:27	1.11


Import	Time Elapsed (min:sec)	Seconds Per Document
1st run	6:52	1.03
2nd run	6:37	1.00
3rd run	6:37	1.00

Image Formats


The Department of Defense uses the CALS standard for exchanging digital forms of technical information. The following Department of Defense documents describe the CALS standard:

- Computer-aided Acquisition and Logistic Support (CALS) Program Implementation Guide (MIL-HDBK-59)
- Automated Interchange of Technical Information (MIL-STD-1840A)
- Raster Graphics Representation in Binary Format (MIL-R-28002).

MIL-STD-1840A describes the standards for two types of documents delivered in digital form: technical publications and product data. Product data consists of engineering and system support data encoded in either of two formats: vector or raster.



DDTS supports only the transfer of product data in raster format. In addition to this limited support of the CALS format, DDTS adds CALSlike support for transfer of FileNet image documents.



CALS Files

A CALS product data document consists of a declaration file followed by a series of data files. Each data file represents an individual image in the document. All DDTS data files are raster data files.

When transferring more than one document to tape in a single transaction, DDTS places all the declaration files together before the data files. The sequence of declaration files and data files for two, two-page documents transferred together is this:

Declaration	Declaration	Raster Data	Raster Data	Raster Data	Raster Data
File D001	File D002	File D001R001	File D001R002	File D002R001	File D002R002

Declaration File

Each document in a transfer is described by a declaration file. The file consists of the identification, source, destination, and classification of a document. The declaration file also gives the count of the files (images) in the document.

The name of the declaration file begins with the letter D followed by a number from 001 to 999. When transferring more than one document in a transaction, the first declaration file is named D001. The numbers increment sequentially for each subsequent document. The declaration file name is an in-transit label that is not used after the transfer completes.

A CALS declaration file consists of sequential, variable-length records. The records are ANSI type D format with a maximum record length of 256 bytes and block lengths of 2048 bytes. The name of the declaration file for each document is placed in the 17-character ANSI file label field.

Declaration File Records

The following table describes the records in each declaration file. The record numbers do not appear in the declaration file; they are included here to show the sequence of records in the declaration file.

Contents of the CALS Declaration File

Record#	Record Identifier	Description
1	SRCSYS:	Source system. Name, address, and other identifying information for the originating system.
2	SRCDOCID:	Source system document identifier. Character string used by the source system to uniquely identify a document.
3	SRCRELID:	Identifier used by the source system to identify a document to which this document is closely related, for example, as a supplement.
4	CHGLVL:	Highest revision and change level in the document. A character string indicating the revision change level and date of this document.
5	DTEISU:	Date of issue of the latest change to the document.
6	DSTSYS:	Destination system. A character string containing the name, address, and other data needed to identify the system to which the document is going.
7	DSTDOCID:	Destination system document identifier. A character string used by the destination system to uniquely identify this document. This is the service or agency document number.
8	DSTRELID:	Character string used by the destination system to identify another document to which this document is closely related, for example, a supplement.
9	DTETRN:	Date of transfer. Date the document was transferred by the source system to the transmission media. The date format is YYYYMMDD. DDTS automatically enters this date during exports.
10	DLVACC:	Delivery account. Free form record giving delivery information specified by the contract.

Contents of the CALS Declaration File, Continued

Record#	Record Identifier	Description
11	FILCNT:	File count. A character string count of the number of each type of data file in the document. Precede each file count with the letter that identifies the type of the data file. DDTS automatically enters this count during export.
12	TTLCLS:	Title security label. A character string stating the security/sensitivity level on the document title.
13	DOCCLS:	Document security label. A character string stating the highest security/ sensitivity level on any file in the document.
14	DOCTYP:	Document type. A character string used by the source system to uniquely identify a document or engineering drawing type, for example, job guide or schematic diagram.
15	DOCTTL:	A character string identifying the document, for example, a technical publication or engineering drawing title.

Sample Declaration File

Following is an example of a declaration file. Each record begins with its record identifier. The identifier string ends with a colon (:) and a space character. Record identifiers are not case sensitive.

```
srcsys: ABC Co., 555 Main St., Costa Mesa, CA 92626
srcdocid: Benchmark 14
srcrelid: Benchmark 12
chalvl: 1
dteisu 19890401
dstsys: ATOS System, Smith Air Force Base, CA 92101
dstdocid: 4SA6-11-4
dstrelid" 4SA6-11
dtetrn: 19900325
dlvacc: Contract Number ABC-12324. Item 4
filant: R12
ttlcls: Unclass
doccls: Unclass
doctyp: Schematic diagram
docttl: QS-78-03 Phase 2
```

Raster Data Files

A raster data file consists of 11 header records followed by image data. The image data in a CALS transfer must be encoded in raster CCITT group 4 code. A FileNet transfer may optionally be encoded in CCITT group 4 code.

All header records are written in the first physical block of the file. The data in the first block is written with a 128-byte ANSI type F, fixedlength record with a 2048-byte block length.

Each raster data file has a three part name (shown in the following table). Like the declaration filename, the raster data filename is an intransit label that is not used after the transfer completes.

Raster Data Filename

Character Positions	Definition
1–4	Name of the declaration file (D001–D999) with which the raster data file is associated
5	Letter R (for "raster")
6–8	3-digit number from 001 to 999

The table on the following page lists the header records in a raster data file. As with the declaration file, the record number does not appear in a raster data file; it merely indicates the sequence of the records in the file.

Each raster header record begins with a record identifier, which is not case-sensitive. The identifier ends with a colon (:) and a space character.

For example, the source system record, SRCDOCID can be written as srcdocid: or SRCDOCID: in the raster data file. All header records are required. When you have no data to place in the record, use the word NONE.

Contents of a Raster Data File

Record #	Record Identifier	Description
1	SRCDOCID	Source system document identifier. An 80-byte string of data elements corresponding to an 80-column card format. Fields in the document identifier are defined in section 5.1.5 of MIL-STD-1840A and section 5.1.9(a) of MIL-STD-804.
2	DSTDOCID	Destination system document identifier. A character string used by the destination system to uniquely identify this document. This is the service or agency document number, identical to Record 7 of the declaration file.
3	TXTFILID	Text file identifier. For a product data file, this record contains the string NONE.
4	FIGID	Figure identifier. For a product data file, this record contains the string NONE.
5	SRCGPH	Source system graphic filename. For a product data file, this record contains the string NONE.
6	DOCCLS	Document security label. A character string stating the highest security/sensitivity level on the data file, identical to Record 13 of the declaration file.
7	RTYPE	Raster data type, where 1 indicates untiled raster graphics data and 2 indicates tiled raster graphics data.
8	RORIENT	Raster image orientation. Two right-justified, 3-character strings separated by a comma, specifying the direction of successive pels (picture elements) and lines. Permissible values are listed in MIL-R-28002.
9	RPELCNT	Raster image pel (picture element) count. Two right-justified, six-character strings separated by a comma, specifying the integer count of pels and lines.
10	RDENSTY	Raster image density. One right-justified, 4-character string representing image density. Permissible values are listed in MIL-R-28002.
11	NOTES	Free form text. The exported image checksum is always stored in 16 bytes at the end of the NOTES record in this format (the underscores represent spaces):_csum:_<8-byte checksum in hex>_

Export

The **ddtsexp** program assembles one or more FileNet documents from storage media into DDTS documents, which are then exported to magnetic disk or tape. The DDTS documents consist of declaration files and the associated raster data files in the formats described in the preceding sections.

Note

For COLD documents with templates, you must export both documents and templates to a compatible system.

The ddtsexp Command

Enter the ddtsexp command as a single line in the following format:

```
ddtsexp [-h] [-e CALS|FILENET] [-i <IMS>] -t | -d <media>
[-c <file1>] [-g] [-l <label>] [-s <file2>] -u <user>
[/<password>] [-p] [-k] [-n] <list>
```

The following table describes the parameters for ddtsexp.

Options for the ddtsexp Command

Option	Explanation
–h	Prints the syntax of the ddtsexp command
-е	CALS specifies a CALS export; FILENET (the default) specifies a FileNet export.
-i	Identifies the NCH name of the Image Services service where the documents to be exported are found (for example, DefaultIMS:pubs:FileNet). If you do not specify an Image Services service, it uses the system default.
-t	Exports to tape (you must be running ddtsexp in the foreground to export to tape). The media name that follows is the NCH name of the drive on which the tape is mounted. You can specify an 8mm tape drive or a 1/2-inch, 9-track tape drive (CALS does not support 1/4-inch cartridge tapes). You can enter a period (.) to specify the default tape drive (which must use CALS-supported densities). You must specify —t or —d, but not both.
-d	Exports to disk. The media name that follows is the directory where you want to place the declaration and data files. You can enter a period (.) to specify the current directory. You must specify –t or –d , but not both.
-с	Creates an import command file during a FileNet export (using the –e option). ddtsimp uses the import command file to import the documents into another FileNet system.
	Use the –c option to specify the name of this import command disk file. When exporting to tape, it creates the file on disk and also copies it to tape. If you do not specify a new filename, the default disk file name is ddtsimp.cmd . The tape file is always named ddtsimp.cmd . The file is not deleted from disk at the end of an export (it may be required for problem resolution); you must delete it manually when you no longer need it.
-g	Exports images in CCITT group 4 format. This option applies only if the –e name is FILENET. CALS images are always exported in CCIT group 4 format.
-I (lowercase L)	Identifies the first four characters of a tape volume label. If you do not specify a label, it uses the letters DDTS the default. The characters you enter are automatically converted to uppercase. (The fifth and sixth characters of the volume label are two digits indicating the number of tapes used.)


Options for the ddtsexp Command, Continued

Option	Explanation
-v	Specifies the maximum number of tape files used for the export. The default (3) is also the maximum for a CALS transfer. The maximum for a FileNet-to-FileNet transfer is 10.
-s	Writes a summary to a file. If you do not specify this option, the summary file goes to stdout.
-u	Identifies the FileNet user whose documents are to be exported. Enter only the object portion of the name. If the user has a password, include a slash (/) and the user's password following the user name. If you enter only the slash, the system prompts you for a password.
-р	Continuously prints to stdout the number of pages transferred
–k	Computes a checksum and compares it for each exported cache object. By default, if checksums don't match, the ddtsexp program aborts. Specify –k to ignore checksum mismatches so that a checksum mismatch merely produces a warning message.
-n	Specifies that the errors generated by the SC_convert tool be treated as non-fatal. When an error is encountered, a warning message will be posted and processing continues on to the next page. The page that has errors will not be written to the output media.
	Lists DDTS export command files. Each command file specifies one or more documents for export.

Export Command Files

Each DDTS export command file starts with the word EXPORT, followed by one or more DDTS document specifications. The global options are FileNet-specific, as are the document option for indexes and the document and page option for annotations.

The following illustration shows the hierarchy of options in an export command file. Loops show iterations. For example, you can specify options for several DDTS documents within one export command file, several FileNet documents within one DDTS document, and several pages within one FileNet document.



Hierarchy of Options Within an Export Command File

Global Options

The ddtsexp program automatically exports each document's ID and class. The document ID, class, and indexes of each DDTS document are those of its first FileNet document. It exports all pages of a FileNet document.

FileNet global options (INDEX and ANNOTATION) export FileNet indexes and annotations and apply only to a FileNet export. You must specify the options in the first file on the file list. You can request a corresponding import command in a DDTS import file for each FileNet global option you specify through an export command. You can modify the import file, if necessary, and use the file to import documents into another FileNet system. See "Page Option" on page 581.

Use the global options described in the following table to apply to all FileNet documents in the export operation. The default for each global option is FALSE.

Option	Explanation
INDEX	Boolean. If INDEX is TRUE, the indexes of all the DDTS documents are exported to the DDTS command file. You can use the indexes as the document indexes on the destination system.
ANNOTATION	Boolean. If ANNOTATION is TRUE, the annotations of all documents are exported. You can import the annotations to the destination system.

Document Options

You can enter values specific to each exported DDTS document. A DDTS document consists of a set of declaration file header records followed by one or more FileNet documents. Each DDTS document section begins with a DOCUMENT command in the following format:

```
DOCUMENT <docname>
```

The <docname> is for tracking only. The identification of a DDTS document is the name of the assigned CALS declaration file. All options affecting a DDTS document, plus the FileNet documents in the DDTS document, are nested within a set of curly brackets {...}. The INDEX and ANNOTATION options are FileNet-specific with a default value of FALSE. The PAGE SRCDOCID and NOTES options refer to values in the raster data files.

The document options are described in the following table.

Option	Explanation
INDEX	Boolean. If INDEX is TRUE, it exports all the DDTS documents' indexes to the DDTS command file. You can use the indexes as the document indexes on the destination system.
ANNOTATION	Boolean. If ANNOTATION is TRUE, it exports all documents' annotations, which you can import to the destination system.
NOTES	Free-form text block, entered in Record 11 of the CALS raster data file
PAGE_SRCDOCID	The source system document identifier. It defines the identifier by a type code, revision, accompanying document type code, distribution statement code, image number, number of images, sheet number and document position code. Refer to section 5.1.5 of MIL-STD-1840A for details on the document identifier and record format, and for tables of applicable codes.

The declaration file records (described in the table, "Contents of the CALS Declaration File" on page 560) follow the document options in the command file. DDTS automatically enters Record 9 (DTETRN) and Record 11 (FILCNT). Do not specify these in the command file.

Page Options

In addition to the global and document options, you can enter values specific to pages in each exported FileNet document (all pages in a FileNet document). Each FileNet document section begins with a FNDOC command in the following format:

```
FNDOC <fndocid>
```

The <fndocid> is the FileNet document number. Enclose FileNet document page numbers and options within the curly brackets. Page numbers and options are nested within the FNDOC command. The pages appear in the following format:

```
PAGE < number >
```

The <number> is the FileNet page number. Enclose within curly brackets the options for that page number. The ANNOTATION option is FileNet-specific. The NOTES and PAGE_SRCDOCID options refer to values in the raster data files. The following table describes these options.

Option	Explanation
ANNOTATION	Boolean. If ANNOTATION is TRUE, the page's annotations are exported and can be imported to the destination system. Setting this value overrides the global and document annotation values.
NOTES	Free-form text block entered in Record 11 of the CALS raster data file
PAGE_SRCDOCID	Source system document identifier, defined by a type code, revision, accompanying document type code, distribution statement code, image number, number of images, sheet number and document position code. Refer to section 5.1.5 of MIL-STD-1840A for details on the document identifier record format and for tables of applicable codes.

Sample Import File

Following is the beginning of a sample DDTS export command file. Compare the format of this file with that of the "Sample Declaration" File" on page 562.

```
EXPORT
DOCUMENT doc1
{INDEX
SRCSYS "ABC Co., 555 Main St., Costa Mesa, CA 92626"
SRCDOCID "Benchmark 14"
SRCRELID "Benchmark 12"
CHGLVL "1"
DTEISU "19890401"
DSTSYS "ATOS System, Smith Air Force Base, CA 92101"
DSTDOCID "4SA6-11-4"
DSTRELID "4SA6-11"
DLVACC "Contract Number ABC-12324, Item 4"
TTLCLS "Unclass"
DOCCLS: "Unclass"
DOCTYP: "Schematic diagram"
DOCTTL "OS-78-03 Phase 2"
# The first DDTS document consists of three FileNet documents.
# Annotations are transferred from page 3 of the first document,
# page 1 of the second document, and page 6 of the third
# document. The source document ID is changed for page 6 of the
# third document.
FNDOC 1234000
{PAGE 3
{ ANNOTATION
FNDOC 1159000
{PAGE 6
{ANNOTATION
PAGE SRCDOCID "Benchmark 13"
```

Import

The **ddtsimp** program imports DDTS documents from tape or magnetic disk and commits them to storage media as FileNet documents. Each imported document can consist of as many as 999 data files, with each data file being an image. Each imported DDTS document is a single FileNet batch with the prefix **d**.

Images and indexes are not verified unless you specifically request verification. On import, if the destination system is configured for checksumming, the system imports images with checksums.

You can commit images to a compatible system (a system with nonoverlapping document ID number ranges) with the same document number, or you can have the importing system assign a new number. Regardless of the DDTS import command, document services assigns new document IDs whenever necessary to avoid duplication.

Note

For COLD documents with templates, you must import both documents and templates to a compatible system.

Indexing information can be imported, or documents can be reindexed. Annotations can be imported or ignored. If you use ddtsimp to import documents from another FileNet system, the documents can be placed into WorkFlo queues.

If an error occurs during import, the batch is gueued in NOQUEUE. Use Document Entry to recover the batch after resolving the cause of failure (check the error log to determine possible causes of failure).

The ddtsimp Command

Enter the ddtsimp command as a single line in the following format:

```
ddtsimp[-h] [-e CALS|FILENET] [-i TIFF/FILENET]
[-b < BES >] -t | -d < media > [-l] [-s | -S] [-f < file >]
[-c <class>] -u <user>[/<password>] [-n] [-k] [-p] list>
```

The following table describes the parameters for ddtsimp.

Options for the ddtsimp Command

Option	Explanation
–h	Prints the syntax of the ddtsimp command
-е	CALS specifies a CALS import. FILENET (the default) specifies a FileNet import.
-i	TIFF imports a document and writes to storage media in TIFF G4 format. FILENET (the default) specifies FileNet G4 format.
-b	Specifies the NCH name of the Batch Entry Service to which the documents are to be imported. If you do not specify - b , the system default is used.
_t	Imports from tape (you must be running ddtsimp in the foreground to export to tape). The media name that follows is the NCH name of the drive on which the tape is mounted. You can use a period (.) to specify the default tape drive.
	You must specify -t or -d, but not both.
-d	Imports from disk. The media name that follows is the directory where the declaration and data files are located. You can use a period (.) to specify the current directory.
	You must specify -t or -d, but not both.
- s	Writes summary information on the DDTS documents specified in the import command files to the file specified in the -f option. No import is performed.
	You can specify -s or -S (or neither), but not both.

Options for the ddtsimp Command, Continued


Option	Explanation
-S	Writes summary information about all DDTS documents to the file specified in the -f option; no import is performed.
	You can specify -s or -S (or neither), but not both.
-f <file></file>	The summary file, specified only with -s or -S . You can enter either a period (.) or stdout to specify the standard output file.
-I (lowercase L)	If you are importing a FILENET file, use this option to load the default ddtsimp.cmd file from tape to disk. You can then examine the import command file and modify it, if necessary, before importing.
	When it encounters an end of file on a reel-to-reel tape, the system looks ahead to see if another file is on the tape. If so, the system opens the file and continues reading. If not, the system rewinds the tape.
	Cartridge tapes cannot backspace, so the system does not look for another file nor rewind the tape. To rewind the tape, enter:
	mt -f <tape-drive-name> rewind</tape-drive-name>
-с	Imports documents using the default document class
-u <user>/<password></password></user>	Specifies the FileNet user who owns the imported documents. Enter only the object portion of the name. If the user has a password, include a slash (/) and the user's password. If you enter only the slash, the systems prompts you for a password.
-n	Do not preserve document IDs even though the transfer is from a compatible FileNet system
-k	Makes checksum errors nonfatal. DDTS issues a warning message but otherwise ignores checksum errors.
-р	Continuously prints to stdout the number of pages transferred
	Lists DDTS import command files (the default is ddtsimp.cmd). Each command file specifies one or more documents for import. If you do not specify an import command file, it imports all documents found in the -d or -t media name, using the default options, into the document class you have specified in the -c option.

Import Command Files

Each DDTS import command file starts with the word IMPORT, followed by one or more DDTS document specifications.

The following figure shows the hierarchy of options in an import command file. Loops show iterations.

For example, you can specify options for several DDTS documents within one import command file, several FileNet documents within one DDTS document, and several pages within one FileNet document.



Hierarchy of Options Within an Import Command File

Global Options

You must specify global options in the first file on the file list. Use the global options described in the following table to apply to all FileNet documents in the import operation.

Option	Explanation
COMMIT	Boolean. If COMMIT is TRUE (the default), an import commits the documents. If COMMIT is FALSE, each DDTS document is put into a FileNet batch and queued in the NOQUEUE.
INDEX	Boolean. If INDEX is TRUE (the default), the system imports the indexes of all the DDTS documents. If INDEX is FALSE, it does not import the indexes and you must index and commit the documents using Document Entry.
VERIFYINDEX	Boolean. If VERIFYINDEX is TRUE (the default is FALSE), an operator must verify all imported indexes. The documents are not committed, even if COMMIT is TRUE.
VERIFYIMAGE	Boolean. If VERIFYIMAGE is TRUE (the default is FALSE), an operator must verify all imported images. The documents are not committed, even if COMMIT is TRUE.
ANNOTATION	Boolean. If ANNOTATION is TRUE (the default), this imports the annotations of all documents. If no document annotation command is given, a global command of FALSE overrides a page-level command. Annotations are created only if the document is committed. Therefore, COMMIT and INDEX must be TRUE and VERIFYINDEX and VERIFYIMAGE must be FALSE.
SKIP	A document number or ALL. Used to restart a partially complete import.
SSN	Identifies the system serial number of the source system that created the declaration and data files. If the source system and destination system are FileNet compatible, it retains the document IDs. If the document IDs are not to be retained, do not specify an SSN (or specify the -n option in your ddtsimp command).
DOCTYPE	Specifies the default type of the documents to be imported: IMAGE, TEXT, FORM, or MIXED. The default for a CALS import is IMAGE. The default for a FILENET import is MIXED.

Document Options

Enter the DOCCLASS command after the global options and before the document options. The DOCCLASS command maps the source system's document class to the destination system's document class, and maps the source system's index names to the destination system's index names. Use this format:

```
DOCCLASS <source_class> IS <destination_class>
{<source index 1> IS <destination index 1>...
<source index n> IS <destination index n>}
```

You can enter values specific to each imported DDTS document. A DDTS document consists of a set of declaration file header records followed by one or more FileNet documents. Each DDTS document section begins with a DOCUMENT command in the following format:

```
DOCUMENT < name>
```

The <name> is the CALS declaration filename. Enter all options pertaining to the DDTS document within a set of curly brackets {...}.

The table on the following page describes these options.

Document Options in the Import Command File

Option	Explanation
DOCID	Specifies the FileNet document ID. Retains the document ID from the source system if the following conditions are true:
	the document ID is specified
	the source and destination systems are compatible
	• the ddtsexp command does not use the -n option
COMMIT	Boolean. If COMMIT is TRUE (the default), the document is committed. If COMMIT is FALSE, the document is put into a FileNet batch and queued in the NOQUEUE. This entry overrides the global commit option for this document.
VERIFYINDEX	Boolean. If VERIFYINDEX is TRUE (the default is FALSE), an operator must verify all indexes imported for the document and manually commit the document even if COMMIT is TRUE. This entry overrides the global commit option for the document.
VERIFYIMAGE	Boolean. If VERIFYIMAGE is TRUE (the default is FALSE), an operator must verify the imported image and manually commit the document even if COMMIT is TRUE. This entry overrides the global commit option for the document.
ANNOTATION	Boolean. If ANNOTATION is TRUE (the default), annotations for the document are imported. This entry overrides the global option. However, if you do not specify a document option, a global option of FALSE overrides the page-level option.
	Since annotations are created only if the document is committed, COMMIT and INDEX must be TRUE and VERIFYINDEX and VERIFYIMAGE must be FALSE.

Document Options in the Import Command File, Continued

Option	Explanation
INDEX	Specifies the index values for the imported FileNet document. The option is in the form:
	INDEX <index_name> = <index_value></index_value></index_name>
	where <index_value> can be a string, a date (yyyy/mm/dd), or an integer. You can specify multiple indexes.</index_value>
DOCTYPE	Specifies the type of the document after import: IMAGE, TEXT, FORM, or MIXED. The default for a CALS import is IMAGE. The default for a FILENET import is MIXED.

Page Option

In addition to the document options, you can use the page option to create annotations for a page. Annotations exported by ddtsexp in hexadecimal format can be imported by ddtsimp.

CAUTION

Do not enter new annotations using ddtsimp, because ddtsimp contains FileNet-specific control information.

The format of the PAGE option is as follows:

```
PAGE < number >
{ ANNOTATION < hex string>
```

The <number> is the FileNet page number. The <hex string> is the annotation placed within either single or double quotation marks (for example, /'hexstring' or /"hex string").

Sample Import File

Following is the beginning of a sample import command file.

```
IMPORT
COMMIT
DOCCLASS enginedrawing IS edrawing
{prod_nos IS pnos
description IS desc
DOCUMENT D001
{DOCID 1234000
DOCCLASS enginedrawing
INDEX prod nos = 1201
}
DOCUMENT D002
{DOCID 1234020
DOCCLASS shopdrawing
}
```

Creating a Script

Instead of typing the ddtsexp and ddtsimp commands each time you transfer documents, consider entering this information into a script. Then modify the script for any changes during a particular run.

When typing the command file, consider duplicating the information for the first document and then changing the document IDs for subsequent documents.

Use an editor such as vi to create a file into which you type the command. For example, you might name the ddtsexp script **export**.

vi export

After saving the file, make it executable using the following command.

chmod a + x export

Finally, move the script to the /fnsw/local/bin directory.

mv export /fnsw/local/bin/export

If you don't have a /fnsw/local/bin directory, create it with the mkdir command.

mkdir /fnsw/local/bin

To run the script without typing the complete path name or changing to that directory every time, make sure /fnsw/local/bin is in your search path. If you need help with this, contact your service representative.

Below is a C shell script that reads a file named doclist in the current directory and generates a DDTS export command file format named **exp.com**. The doclist file contains document IDs that you want to export.

```
#! /bin/csh -f
# This script reads the file "./doclist" which contains document
# IDs and generates an output file "./exp.com" in DDTS (Tape
# Transfer) export command file format. The output file can then be
# edited prior to DDTS exports.
# usage: createcom
set infl = "./doclist"
set outfl = "./exp.com"
set counter = 1
echo "EXPORT" > $outfl
echo "INDEX TRUE" >> $outfl
echo "ANNOTATION TRUE" >> Soutfl
foreach f ('cat $infl')
   echo "-----"
   echo "Processing document $f"
   echo "-----"
   echo "" >> Soutfl
   echo "DOCUMENT doc$counter" >> $outfl
   echo "{" >> $outfl
   echo "FNDOC $f" >> $outfl
   echo "}" >> $outfl
   @ counter=$counter + 1
end
echo "Output file is " $outfl
echo "End of the script."
```

For example, if doclist contains

100001 100002

and you run the script, createcom, it creates the file exp.com which looks like the following.

```
EXPORT
INDEX TRUE
ANNOTATION TRUE
DOCUMENT doc1
FNDOC 100001
DOCUMENT doc2
FNDOC 100002
```

You can then edit the file (or the script) to meet your criteria.

Appendix A – Function Codes

The following tables list the function codes.

Check spelling and capitalization before saving the entry. The program does not validate your entry. Application Executive

Function Codes for Application Executive

Menu Option	Function Code	Function Name
Database Maintenance	dbmaint	Database Maintenance
Security Administration	sysadmin	Security Administration
Storage Library Control	servercontrol	Storage Library Control
Background Job Control	scpbkg	Background Job Control
Cache Export/Import	cache_export_import	Cache Export/Import

Server Print Program

Function Code for Server Print Program

Menu Option	Function Code	Function Name
Server Print Program	printing	Server Print Program

Database Maintenance

Function Codes for Database Maintenance

Menu Option	Function Code	Function Name
Archiving Utilities	arcutil	Archiving Utilities
Delete Doc./Folder	deldoc	Delete Doc./Folder
Update Doc. Security	updatedocsec	Update Doc. Security
Update Retention Parameters	updateretent	Update Retention Parameters
Define/Update Index	defineindex	Define/Update Index
Rename Index	renameix	Rename Index
Build Retrieval Key	buildix	Build Retrieval Key
Drop Retrieval Key	dropix	Drop Retrieval Key
Define Media Family	defineoptical	Define/Update Family
Define/Update Cluster	definecluster	Define/Update Cluster
Index Report	indexreport	Index Report
Build Menu	custbuild	Build Menu
Define/Update Class	definedocclass	Define/Update Class
Class Report	classreport	Class Report
Define/Update Family	defineoptical	Define/Update Family
Family Report	opticalreport	Family Report

Storage Library Control

Function Codes for Storage Library Control

Menu Option	Function Code	Function Name
Main screen	servercontrol	Main screen
Detailed Surface Info	disksurface	Detailed Surface Info
Pending Surface Requests	diskdetail	Pending Surface Requests
Media Space Usage	diskusage	Media Space Usage
Local/Foreign ID's	dispsurfids	Local/Foreign IDs
Create Doc Header File	makedhfile	Create Doc Header File
Eject Media by Surface ID	ejectoptical	Eject Media
Disable/Enable Media	enableosar	Enable/Disable Storage Media
Change Media Type	changedisk	Change Media Type
Media Family Info	dispfaminfo	Media Family Info
Change Family Name	changefam	Change Family Name
Local Statistics	dsusage	Local Statistics
Remote Committals	rmtcomreq	Remote Committals
Enable/Disable Library	enableosar	Enable/Disable Storage Library
Respond to RSVP	replyrsvp	Respond to RSVP
Delete Info Message	deleteinfo	Delete Info Message
Configure Library screen	osarstatus	Configure Library screen
Insert Media	acceptoptical	Insert Media
Eject Media	ejectoptical	Eject Media
Preformat Media	formatdisk	Preformat Media
Slot Drive Map	displaymap	Slot Drive Map
Media Surface Summary	osardetail	Media Surface Summary
Media Space Usage	diskusage	Media Space Usage

	Function	Codes for	Storage	Library	Control.	Continued
--	----------	-----------	---------	---------	----------	-----------

Menu Option	Function Code	Function Name
Enable/Disable Slot	enableosar	Enable/Disable Storage Library
Enable/Disable Grippers	enableosar	Enable/Disable Storage Library
Calibrate Library	calibrateosar	Calibrate Library
Identify Media in Library	maposar	Identify Media in Library
Enable/Disable Drive	enableosar	Enable/Disable Storage Library
MSAR Backup Mode	msarbackup	MSAR Backup

Background Job Control

Function Codes for Background Job Control

Menu Option	Function Code	Function Name
Incorporate Foreign Media	insertforeign	Incorporate Foreign Media
Manually Incorporate Foreign Media	insertmanual	Manually Incorporate Foreign Media
Copy Documents	copydisk	Copy Documents
Copy Documents Using File	copywfldocs	Copy Documents Using File
Copy Annotations From Database to Media	copyanno	Copy Annotations From Database to Media
Consolidate Media	consoldisk	Consolidate Media
Erase Media	erasemedia	Erase Media
Rebuild Media	rebuilddisk	Rebuild Media
Import Documents From Media	importdisk	Import Documents From Media
Create Archive Database	createarchive	Create Archive Database

Function Codes for Background Job Control, Continued

Menu Option	Function Code	Function Name
Import Archive Image Service	importarchims	Import Archive Image Service
Find Open Documents	findopendocs	Find Open Documents
Completed Jobs	dispbkgerrors	Completed Jobs
Results of Find Open Documents	dispopendocs	Results of Find Open Documents
Modify Status of Background Job	modbkgstatus	Modify Status of Background Job
Delete Log of Completed Job	compjobdellog	Delete Log of Completed Job
Convert optical surface to MSAR	msarconvert	MSAR Convert

Cache Export/Import Program

Function Codes for Cache Export/Import Program

Menu Option	Function Code	Function Name
Export Cache Objects	backupobjects	Export Cache Objects
Show Cache Objects	showobjects	Show Cache Objects
Import Cache Objects	restoreobjects	Import Cache Objects

COLD Application

Function Codes for COLD Application

Feature	Function Code	Function Name	
Define Background Template	coldIItemplate	Define Background Template	
Define Channel Control File	coldlichannel	Define Channel Control File	
Define Report Format	coldIIreportformat	Define Report Format	
Define Import Job	coldllimportjob	Define Import Job	
Preview Documents	coldIlpreview	Preview Documents	
Import Documents	coldIlimport	Import Documents	
View Import Log	coldlimportlog	View Import Log	

Overriding a Busy Batch

Function Code for Overriding a Busy Batch

Special Function	Function Code
Overriding a Busy Batch	DE_BUSYOVERRIDE
PC workstation users can only override busy batches if they belong to the SysAdminG or DE BUSYOVERRIDEG groups.	

Full Functions Names Listing

Full Function Names Listing

Function Name	Function Code
Archiving Utilities	arcutil
Background Job Control	scpbkg
Build Menu	custbuild
Build Retrieval Key	buildix
COLD Main Menu	coldII
Cache Backup Program	cachebackup
Cache Export/Import	cache_export_import
Calibrate Library	calibrateosar
Change Family Name	changefam
Change Media Type	changedisk
Class Report	classreport
Completed Jobs	dispbkgerrors
Configure Library Screen	osarstatus
Consolidate Media	consoldisk
Copy Annotations From Database to Media	copyanno
Copy Documents	copydisk
Copy Documents Using File	copywfldocs
Create Archive Database	createarchive
Create Doc Header File	makedhfile
Database Maintenance	dbmaint
Define Background Template	coldlitemplate
Define Channel Control File	coldlichannel

Full Function Names Listing, Continued

Function Name	Function Code
Define Import Job	coldlimportjob
Define Report Format	coldIlreportformat
Define/Update Class	definedocclass
Define/Update Cluster	definecluster
Define/Update Family	defineoptical
Define/Update Index	defineindex
Delete Doc./Folder	deldoc
Delete Info Message	deleteinfo
Delete Log of Completed Job	compjobdellog
Delete RSVP	deletersvp
Detailed Surface Info	disksurface
Drop Retrieval Key	dropix
Eject Media	ejectoptical
Eject Media by Location	ejectbyloc
Enable/Disable Storage Library	enableosar
Erase Media	erasemedia
Export Cache Objects	backupobjects
Family Report	opticalreport
Find Open Documents	findopendocs
Identify Media in Library	maposar
Import Archive Image Service	importarchims
Import Cache Objects	restoreobjects
Import Documents	coldlimport
Import Documents From Media	importdisk

Full Function Names Listing, Continued

Function Name	Function Code
Incorporate Foreign Media	insertforeign
Index Report	indexreport
Insert Media	acceptoptical
Local Statistics	dsusage
Local/Foreign IDs	dispsurfids
MSAR Backup	msarbackup
MSAR Convert	msarconvert
Manually Incorporate Foreign Media	insertmanual
Media Family Info	dispfaminfo
Media Family Information	dispwritesurf
Media Space Usage	diskusage
Media Surface Summary	osardetail
Migrate Documents	migratedocs
Modify Status of Background Job	modbkgstatus
Pending Surface Requests	diskdetail
Preformat Media	formatdisk
Preview Documents	coldIlpreview
Rebuild Media	rebuilddisk
Remote Committals	rmtcomreq
Rename Index	renameix
Respond to RSVP	replyrsvp
Results of Find Open Documents	dispopendocs
Security Administration	sysadmin
Server Print Program	printing

Full Function Names Listing, Continued

Function Name	Function Code
Show Cache Objects	showobjects
Slot Drive Map	displaymap
Storage Library Control	servercontrol
Update Doc. Security	updatedocsec
Update Retention Parameters	updateretent
View Import Log	coldIlimportlog

Appendix B – Date and Time Formats

This appendix describes the available date and time formats and how the FileNet Image Services server applies its date mask functionality. It identifies the supported Image Services date and time formats and notes exceptions to the standard implementation.

Note

The Image Services software supports the year 2000 standard date requirement as specified in the ANSI 3.30 and ISO 8601 standards documents.

Current Date Mask Functionality

This section describes the FileNet date mask functionality for the Image Services software described in this manual. This functionality applies to both the new and old date mask applications in the 3.5.0 release of the Image Services software, except where specifically noted.

For details, see the following topics:

- "Support for Local Masks" on page 597
- "Date and Time OS Settings" on page 597
- "Log File Date Format Exceptions" on page 602
- "New Default Date Mask Application" on page 603
- "Old Default Date Mask Application" on page 604

- "Conversion of 2-Digit to 4-Digit Date Format" on page 607
- "Date and Time Number Conversion Ranges" on page 608

Support for Local Masks

The Image Services software provides new default date mask functionality to support both a 4-digit date format and the international standard date format based on the server's platform configuration settings.

The date and time formats vary, depending on the common usage in a particular locale. For example, while a server installed in Europe might use the dd/mm/yyyy format, a server installed in the United States most likely uses the mm/dd/yyyy format.

Date and Time OS Settings

The method used to set the date and time through the server's operating system settings is platform specific:

- A UNIX system determines the date and time masks through the current locale. (For configuration details, see the UNIX man page entry for "locale". For SUN configuration procedures, see "SUN Platform Date and Time Settings" on page 598.)
- A Windows Server system uses the date and time masks configured through the Control Panel's Regional Settings panel. (For configuration details, see the related Windows Server on-line help system.)

The Image Services software obtains the default date and time masks only once per process. Changes made to these settings affect the Image Services system defaults only after restarting the application.

Note

You should configure your operating system with a 4-digit year date mask.

SUN Platform Date and Time Settings



To configure the date and time mask on a SUN platform, you must compile a locale configuration file on a system with a compiler resident, then copy the files to the correct directory path on the Image Services server. Follow these steps:

- Obtain the locale configuration source file and compiler from SUN Mi-1 crosystems.
- 2 Install the compiler onto the SUN system you want to use for compiling the locale configuration source file. If you want to use a different system than the Image Services server, make sure the OS software release matches the one installed on the Image Services server where you plan to copy the compiled source file.
- 3 Log on as root to the SUN platform where you installed the compiler and copy the locale source file into a working directory.
- Copy the source file from the working directory into another directory. For example:

cp <any directory>/<source from SUN>.src <any directory>/ <new locale>.src

5 Use an editor, such as vi, to edit the <new locale>.src file to change the date and time masks shown below. The following table identifies the mask type, keyword, and sample values.

Mask Type	Keyword	Example
date	d_fmt	%Y-%d-%m for yyyy-dd-mm
time	t_fmt	%I:%M:%S %p for hhhh:mm:ss
date and time	date_fmt	%Y-%d-%m for yyyy-dd-mm %I:%M:%S %p for hhhh:mm:ss

- 6 Compile the source and copy to the appropriate directory as shown for your OS version:
 - For Solaris 2.5.1:
 - Enter the command:

localedef -c -f <charmap> -i <new_locale>.src <new_ locale>

For example:

localedef -c -f charmap.iso 2.5.1 -i dates.src dates

The compile program generates a few files and attaches a .time suffix to the <new locale> compiled file name. For the example above, the compiled file name is **dates.time**.

Note

If you cannot obtain the charmap file from Sun, edit the .src, delete all sections except for LC_TIME, then run the compile without the charmap option (for example: localedef -c -i dates.src dates).

b Rename the compiled file (dates.time in the example above) to time and copy it to the Image Services server under the following directory:

/usr/lib/locale/<locale dir>/LC TIME

where <locale dir> is the directory that identifies your locale, such as **en_UK** for the United Kingdom locale environment.

- For Solaris 2.6:
 - Enter the command:

localedef -c -i <new locale>.src <new locale>

For example:

localedef -c -i dates.src dates

The compile program generates two files.

b Rename the file in the form <newlocale>.so.1 to the following directory:

/usr/lib/locale/<locale dir>

where <locale dir> is the directory that identifies your locale, such as **en UK** for the United Kingdom locale environment.

- 7 Make sure that bin is the group and owner of the time file and set it to bin if it is not.
- 8 Change the permission on time to 555 (read and execute for all users).

- Set the locale as shown for your OS version: 9
 - For Solaris 2.5.1:

Set the directory using the your shell environment command. For example, if you copied the edited time file to directory path for en UK, enter one of these commands:

```
setenv LC TIME en UK (C shell)
export LC TIME=en UK (K shell)
```

Note

To revert this back to the default directory, enter one of these commands: **setenv LC_TIME C** (C shell) or **export LC_TIME=C** (K shell).

For Solaris 2.6:

Set the directory using the your shell environment command. For example, if you copied the edited time file to directory path for en_ UK, enter one of these commands:

```
setenv LC ALL en UK (C shell)
export LC ALL=en UK (K shell)
```

Note

To revert this back to the default directory, enter one of these commands: setenv LC ALL " " (C shell) or export LC ALL=" " (K shell).

10 Test the new date format by entering the command:

date

The system displays the date using the modified format.

Note

If you want to set the directory permanently for all users who login to the Image Services server, set the locale directory in the /etc/.login file as described for your shell environment above.

Log File Date Format Exceptions

In general, log files displayed through a graphical user interface (GUI) Image Services application use the OS-configured date format. Changes made to the default date mask affect any process that starts after the change.

Although the Image Services system forces the user to enter a 4-digit year, Image Services displays the log file date using a 2-digit format for all log file entries recorded when the operating system is set with a 2digit year date mask. Any log files created before you changed a date mask and restarted the application retain the previous mask.

While most Image Services GUI applications apply the default date format, the following applications use hard-coded entry dates:

- Task Manager (Xtaskman) displays the log list (available through the Monitor menu's Event Log option) using the date format of yyyy/mm/dd. For example, all entries to events that occurred on September 8, 1999 show the date as 1999/09/08.
- Syslog (elog) entry dates display the date format as part of the event log file name using the format of yyyymmdd. The Image Services system stores these files in the event log directory.

For example, you can find an event log file generated on September 8, 1999 under the default file path shown for your operating system:

/fnsw/local/logs/elogs/elog19990908 for UNIX servers

<drive>:\fnsw loc\logs\elogs\el19990908 for Windows Servers

Syslog files contain entry dates using the date format of yyyy/mm/ dd. For example, error messages stored in elog19990908 on a UNIX platform may list messages with the date "1999/09/08."

Wizard log files use the date and time format of YYYY/MM/DD HH:MM:SS. For example, when an installation wizard creates a log file of its progress, it records an event that occurred at 1:30 PM on November 20, 1999 as 1999/11/20 13:30:00.

New Default Date Mask Application

The Image Services system supports the 4-digit year requirement by applying the date mask as follows:

- Where the Image Services system prompts you with a yyyy format, it does not accept less than a 4-digit date. For example:
 - If you enter 97 or 097, Image Services refuses to accept it and displays an error message.
 - If you enter 0097, Image Services accepts it as the year 97, not 1997.
- If the operating system is configured with a 2-digit year date mask, the Image Services system converts it to a 4-digit year for data entry.

Old Default Date Mask Application

Existing customers may have legacy applications that require the previous (old) date mask functionality to access FileNet records. If your application requires the Image Services system's previous functionality for date formats, you must set the MASK OLD option as described in the following subsections for your operating system.

CAUTION

You should avoid reverting the date mask functionality to its previous method since it could cause undesirable date conflicts.

The Image Services system's old functionality handles date masks as follows:

- Where the Image Services system prompts you for a yyyy date, Image Services accepts the date entered without applying any changes. For example, if you enter 97, 097, or 0097, Image Services accepts it as the year 97, not 1997.
- Where the Image Services system prompts you for a yy date and you enter more than two digits (such as 1997), Image Services refuses to accept it and displays an error message.

To preserve the previous date mask functionality, you must set the MASK_OLD option to 1, as described for the platform on which the Image Services software resides. If MASK OLD is not present, the Image Services system uses the new mask functionality.

Preserving the Old Date Mask on a Windows Server System



To preserve the old functionality on a Windows Server-based Image Services server, you must add a registry entry for MASK OLD.

Note

Before you can edit the Windows Server system registry, you must log onto the Windows Server system using a user account that has permission to modify the registery.

To add the MASK_OLD setting in the Windows Server registry, follow these steps:

1 Run **regedt32** and navigate to the following key under HKEY LOCAL MACHINE:

SOFTWARE\\FileNet\\<application>\\CurrentVersion

where:

<application> is either: ISTK on the client machine IMS on the Image Services Windows Server

- 2 Select the appropriate key described above.
- 3 Open the Edit menu and select the Add Value option.
- Type in MASK_OLD in the Value Name field. 4
- 5 Select **REG_SZ** for the Data Type.
- 6 Enter a single 1 as the String.

7 Save your changes and restart the Image Services server to activate the changes.

Preserving the Old Date Mask on a UNIX System



To preserve the old functionality on a UNIX-based Image Services server, you must set the MASK OLD environment variable to 1.

Note

Before adding an environment variable to the appropriate file, you must log onto the UNIX system using a user account that has permission to modify the file that sets the environment variables.

To set the MASK OLD environment variable to 1, follow these steps:

1 Using a text editor, add the appropriate environment variable in the .profile or .cshrc file for the users who require the old mask functionality, as shown for your shell type:

> csh: setenv MASK OLD 1 bsh or ksh: export MASK OLD=1

2 Save your changes and restart the Image Services server to activate the changes.

Note

This procedure only affects the UNIX system environment variable for the user defined in the login script to which you're adding this command. Consult the OS system administrator to determine which file the operating system uses to set the environment for all users when the Image Services system starts automatically upon system startup.

Conversion of 2-Digit to 4-Digit Date Format

The Image Services system (including the server-side Image Services and client-side ISTK software) requires the user to enter a 4-digit year. If the operating system date mask is set with a 2-digit year, the Image Services system converts it to a 4-digit year when prompting the user to enter a date.

The following example shows a date mask used to prompt the user for a document entry date.



In the above example, the UNIX system on which this Image Services system resides has been configured with a date mask of mm/dd/yy. Image Services has converted the date mask to display a 4-digit year format. The user must enter the date using the displayed format.

Reports related to dates for documents defined on an operating system configured with a 2-digit year mask display the actual dates using the 2-digit format. If the user enters the date of 01/02/1999 in place of the date mask shown above, a report accessing records with these document entry dates shows the date as 01/02/99 since it applies the date mask set in the operating system.

Each application (including user-defined APIs) that calls the date and time from the Image Services server determines whether to use a 2digit year format or the standard 4-digit year format. If an application displays a 2-digit year date mask in user prompts, the Image Services system accepts either 2-digit or 4-digit year entries. The Image Services system appends the current century to 2-digit year entries.

For example, when presented with a date mask of yy/mm/dd, a user may enter the date as 99/01/02.

- If the user enters this date when the server is set to the year 1999, Image Services converts the date to 1999/01/02.
- If the user enters this date when the server is set to the year 2000, Image Services converts the date to 2099/01/02.

Date and Time Number Conversion Ranges

The Image Services software converts and stores date entries as numbers based on the year 1970. This method ensures accurate date conversions, whether the original date entries were made using a 2-digit or 4-digit year date mask.

The Image Services software converts dates differently, based on whether the date mask includes time:

- If the date mask does not include the time, the Image Services system converts each date entered to the number of days before or after January 1, 1970.
- If the date mask includes the time, the Image Services system converts each date and time entered to the number of seconds before or after January 1, 1970.

The converted number for date can only be 32-bits in size. This size requirement does not have any practical limits for numbers converted from date-only masks. However, this restriction limits the converted numbers to dates within an approximate range between 1906/08/17 and 2038/01/09 when using a date with time mask.

Supported Date and Time Formats

The Image Services software supports the OS-configured date and time formats, as described in the following topics:

- "Valid Separators" on page 610
- "Default Formats" on page 612
- "Date Formats" on page 614
- "Time Formats" on page 615
- "Windows Server Date and Time Format Conversion" on page 616

Valid Separators

The Image Services system supports the following date and time separators:

```
\{ \text{actual space} > , . / ? : ; ````() < > [] { } -= + | * & ^ % $ # @ !
```

The operating system may not support all FileNet-supported separators, as described below.

OS Exceptions

Some of the characters acceptable as FileNet separators have a special meaning to the operating system. Although you could use these characters internally in the Image Services software as mask separators, they may produce errors or incorrect values when a user enters a date or the system displays the mask.

To prevent potential problems, you should avoid using the following characters as separators:

- Double (") or single (') quote characters: The UNIX and DOS shells use these characters to designate the start and end of a string. When you enter a date mask surrounded by quotation marks at a command line, the system fails to recognize the format and generates and error.
- Percent sign (%) character: The C library printf() function uses this character. If you have defined a mask with a percent sign, the system may not display the mask or date properly and the program may generate an error.
- OS-specific special characters: The UNIX shell and the DOS command window assign special meaning to many characters. Using

any of these characters in the mask may cause problems when an operator enters the mask or date on a command line.

However, you may still use most of these special characters on the command line by preceding them with an escape code, such as the backslash (\) or by surrounding the entire mask or date with the single quote mark, such as 'mm\$dd\$yyyy'.

Masks without Separators

You may also specify masks without separators. (For example, yymmdd is valid.)

However, when you do not use separators, the date entered must not be ambiguous. To avoid ambiguity, you must enter the date as two digits when using single digit masks such as h, m, or d. The user must enter two digits into all fields displayed with a mask that can accept two digits, even though the mask indicates a single digit. The number of characters entered must match the maximum possible number for the mask unless the final number is a single digit (0-9).

For example, yyyymd is a valid mask. However, a user should enter a date of January 2, 1999 as 19990102 not 199912.

- If the user enters the date as 1999012, Image Services accepts it as January 2, 1999.
- If the user enters the date as 1999102, Image Services accepts it as October 2, 1999.

Default Formats

Image Services replaces unsupported date and time masks (configured in the operating system) with default standard masks for use with FileNet records.

Note

The Task Manager (Xtaskman), initfnsw, and whatsup do not depend on any FileNet shared libraries, so they do not replace unsupported date and time masks with the default masks. Instead, they rely on the system API to format the date and time.

Windows Server Default Masks



When encountering unsupported date or time masks on a Windows Server platform, Image Services replaces the mask with the following default masks:

Туре	FileNet Mask	U/I Display Mask*
Date	yyyy/mm/dd	yyyy/MM/dd
Time	HH:tt:ss	HH:mm:ss

^{*} The mask displayed through the application's user interface applies the Windows Server-specific mask, rather than the default FileNet mask.

Note

If the Task Manager, initfnsw, or whatsup encounters an unsupported date or time mask on an Windows Server platform, the Windows Server system replaces the mask with the standard date mask of yyyy/ mm/dd with a 24-hour clock.

UNIX Default Masks



When encountering unsupported date or time masks on a UNIX platform, Image Services replaces the mask with the following default masks:

Туре	FileNet Mask	U/I Display Mask*
Date	yyyy/mm/dd	yyyy/mm/dd
Time	HH:tt:ss	HH:tt:ss

^{*} The mask displayed through the application's user interface matches the default FileNet mask.

The Image Services system does not display an error message when it encounters an unsupported date mask. Instead, Image Services records the event in the elog where it identifies the error and default date mask.

Date Formats

The Image Services server supports date formats as described in the following table.

Туре	Mask	Description
Day	d*	Numerical day of month without a leading zero
	dd	Numerical day of month with a leading zero for 1-digit numbers
	day	Abbreviation for name of weekday (example: Mon)
	dayname	Full name of weekday (example: Monday)
	ddd	Numerical day of year (1-366)
	W	Numerical day of week (Windows Server: 0=Mon UNIX: 0=Sun)
Month	m*	Numerical month of year without a leading zero
	mm	Numerical month of year with a leading zero for 1-digit numbers
	mon	Abbreviation for name of the month (example: Jan)
	month	Full name of the month (example: January)
Year	уу**	2-digit year
	уууу	4-digit year

- * If you use a single-digit mask, you must either use separators or 2-digit entries when entering the full date. Image Services generates an error when trying to parse ambiguous entry dates.
- ** The new date functionality forces a 4-digit year mask (yyyy) for data entry, even when the operating system mask is set for a 2-digit year (yy). However, each subsystem may apply a 2-digit year mask in its screen displays or generated reports.

Time Formats

The Image Services server supports time formats, as described in the following table.

Туре	Mask	Description
Hour	h*	Hour of day without a leading zero (12 or 24 hour time clock)
	hh	Hour of day with a leading zero for 1-digit numbers (12 or 24 hour time clock)
	H*	Hour of day without a leading zero (forces a 24-hour time clock)
	НН	Hour of day with a leading zero for 1-digit numbers (forces a 24-hour time clock)
Minute	t*	Number of minutes within the hour without a leading zero
	tt	Number of minutes within the hour with a leading zero for 1-digit numbers
Second	s*	Number of seconds within the minute without a leading zero
	SS	Number of seconds within the minute with a leading zero for 1-digit numbers
AM/PM	am**	Used with a 12-hour clock: am or pm

- * If you use a single-digit mask, you must either use separators or 2-digit entries when entering the time. Image Services generates an error when trying to parse ambiguous entry times.
- ** Used to indicate a 12-hour clock only when used with a lower-case hour mask. An upper-case hour mask overrides this mask. The Image Services system uses am, while the Windows Server operating system uses tt for this mask.

The following example shows how different time masks configure the Image Services system to display 1:02:03 PM:

hh:tt:ss am = 01:02:03 PM 13:02:03 hh:tt:ss HH:tt:ss am =13:02:03 PM HH:tt:ss = 13:02:03

Windows Server Date and Time Format Conversion

The Image Services system provides a uniform method of retaining date and time formats, while still allowing for platform-specific date and time masks on a Windows Server system.

The Image Services system enables Image Services subsystems to display Windows Server system-configured date and time formats through the subsystem's user interface or in generated reports. To maintain compatibility of data retention values, the Image Services system converts and stores dates and times using FileNet-standard masks that are compatible with a UNIX operating system.

CAUTION

Do not try to use a literal text mask. Image Services does not support the Windows Server feature of inserting literal text into a mask by surrounding it with single quotes. Image Services interprets the characters within the quotes as a mask. This may cause a variety of unpredictable date and time mask errors.

The date and time masks configured through the Windows Server Control Panel's Regional Settings window do not comply with FileNet standards. The Image Services system converts the incompatible masks for FileNet use. (See "Default Formats" on page 612.)

The Image Services system converts and stores the Windows Serverspecific date and time masks using FileNet standard masks, as shown in the following table.

Ma	ask	
Windows Server	FileNet	Description
dddd	dayname	Full name of weekday
М	m	Numerical month of year without a leading zero
MM	mm	Numerical month of year with a leading zero for 1-digit numbers
MMM	mon	Abbreviation for name of month
MMMM	month	Full name of month
m	t	Number of minutes within the hour without a leading zero
mm	tt	Number of minutes within the hour with a leading zero for 1-digit numbers
tt*	am	Used with a 12-hour clock: am or pm

^{*} Some locales (such as the UK) do not use any indicator for am or pm, which causes problems when using the tt mask. To avoid conflict with locale standards, do not use the tt mask. However, if you must use the tt mask, you must also insert values for the AM and PM symbols in the Regional Settings Time menus. (For configuration details, see the related Windows Server on-line help system.).

The following conversions are not reversible because they are not unique.

	Mask		
Windows Server	FileNet	GUI	Description
у	уу	уу	2-digit year
ууу	уууу	уууу	4-digit year
t	am	tt	Used with a 12-hour clock: am or pm

Note

Both Windows Server and Xtaskman use the ddd mask as an abbreviated name of weekday. All other Image Services functions use ddd as the Julian day of year (1-366). Image Services does not convert ddd to its Image Services equivalent day.

Appendix C – Logic for Retrieving Surfaces and **Ejecting Media**

This appendix is for use as a technical reference and describes Image Services software logic.

Logic for Retrieving Surfaces

The Image Services software uses the following logic for choosing a primary or tranlog surface for retrieval.

Note

This logic remains unchanged for MSAR support with one exception: When both MSAR and optical surfaces are available, the MSAR surface will be favored. The logic continues until the requested surface has been located.

- 1 If one surface is available and another surface is not available, Image Services selects the surface that is available. IS uses the following criteria to determine if a surface is not available:
- 2 The surface is not in a library, "out of box", then:
 - The surface SRF "Do-Not-Use" flag is set to TRUE.
 - b The surface is in a slot and the slot is disabled. (On an MSAR library disabling of a slot is not supported.)
 - c The surface is in a slot and the gripper is disabled. (On an MSAR library disabling of a gripper is not supported.)

- The surface is in a slot and all the drives in the library are disabled.
- e The surface is in a drive and the drive is disabled. (This does not occur for MSAR libraries because when a drive is disabled the surface in the drive is always moved to a slot. This may occur on optical libraries because a drive may be disabled and the gripper may have been disabled and as a result the surface is in a disabled drive.)
- f The surface is an optical two-sided surface, loaded on the other side in a drive and the gripper is disabled.
- The surface is an optical two-sided surface, loaded on the other g side in an ODU (Optical Disk Unit, stand alone drive with no gripper) library.
- h The surface is in a gripper and the gripper is disabled. (On an MSAR library disabling of a gripper is not supported.)
- 3 If both surfaces are "out of box" and:
 - There are no outstanding requests for either surface, select the primary surface.
 - There are outstanding requests for the primary surface and none for the tranlog surface, select the primary surface.
 - c There are outstanding requests for the tranlog surface and none for the primary surface, select the tranlog surface.
 - There are outstanding requests for both surfaces, select the surface with more high priority read requests.

e There are outstanding requests for both surfaces and the surfaces have the same number of high priority read requests, select the primary surface.

Note

The above algorithm is applied with no distinction between MSAR and optical surfaces unless noted.

- Both surfaces are in the box: 4
 - If one surface is an MSAR surface and the other surface is an optical surface, select the MSAR surface.
 - If either of the surfaces are two-sided optical disk, make the following selection:
 - If one surface is a one-sided surface and one surface is a twosided surface that has queued requests on the other side, select the one sided surface.
 - If both surfaces are two-sided surfaces and the primary surface does not have requests on the other side and the tranlog surface has requests on the other side, select the primary surface.
 - If both surfaces are two-sided surfaces and the tranlog surface does not have requests on the other side and the primary surface has requests on the other side, select the tranlog surface.

This is done to reduce flipping of two-sided media. If none of the rules above apply, move on to the next selection criteria.

If one surface is in a drive and the other surface is not, select the surface that is in the drive.

- d If both surfaces are in drives, select the surface with the fewer high priority read requests.
- e If both surfaces are in drives and the number of high priority read requests are the same, select the primary surface.
- If both surfaces are not in drives but both surfaces are "in the box" (i.e. both are in slots), and the primary surface has no outstanding requests, and the tranlog surface has outstanding requests, select the tranlog surface.
- If both surfaces are not in drives but both surfaces are "in the box" (i.e. both are in slots), and the tranlog surface has no outstanding requests, and primary surface has outstanding requests, select the primary surface.
- h If either of the surfaces is a two-sided optical disk, make the following selection at this point:
 - If both surfaces are loaded into a drive on the other side, select the tranlog surface.
 - If one surface is loaded into a drive on the other side, select the other surface.

This is done to avoid selecting upside down media to queue requests to and as a result to reduce flipping of two sided media that are in drives.

If both surfaces are not in drives but both surfaces are "in the box" (i.e. both are in slots), select the surface with more high priority read requests.

If both surfaces are not in drives but both surfaces are "in the box" (i.e. both are in slots) and the surfaces have the same number of high priority read requests, select the primary surface.

Logic for Ejecting Media

When a library becomes full and it's necessary to eject a surface, Image Services follows this logic.

- 1 Find the slot that contains the labeled surface with the oldest mount time in the library. Only select a surface that has no requests. For optical, if the disk in question is a two-sided disk, select a disk with no request on both sides.
- 2 If no slot is found that meets the above criteria, select a slot that contains an unknown or unlabeled blank surface to eject. This only applies to optical libraries.
- 3 If no slot is found that meets the above criteria, select a slot that contains a surface with the oldest mount time with requests. An unlabeled MSAR surface will not be ejected.
- 4 For an optical library, look in the drives for a surface to eject.

Appendix D – Task Manager Configuration File

By default, Image Services permits users in the fn admin and fn op group s to remotely start, stop or query the status of any host running Image Services that is reachable across the network. Remote access capabilities are helpful to many administrators and many backup and restore scripts rely on them. However, they can also be seen as a security risk, particularly at those sites with multiple, widely distributed servers.

Note

The Task Manager Configuration file is optional. If it is not present, all programs will default to the old, non-restricted behavior.

The Task Manager Configuration file gives administrators the ability to restrict remote access. It currently supports two options:

restrict outgoing: Disables "-h HOSTNAME" (vl., initfnsw, whatsup) and [Connect] (Xtaskman).

restrict incoming: TM daemon refuses all remote requests from other servers.

TM daemon (Task Manager service), Xtaskman (FileNet Task Manager), initfnsw, whatsup and vI all support this configuration file.

The file is located in the following directory (depending on platform):

/fnsw/local/sd/1/tm config.txt



<drive:>\\fnsw loc\\sd\\1\\tm config.txt

Sample Configuration File

```
tm config.txt: Configure TM daemon and ev tools
 LOCATION: <drive:>\\fnsw_loc\\sd\\1\\tm_config.txt (Windows)
                   /fnsw/local/sd/1/tm config.txt (UNIX)
 OPTION
                        DESCRIPTION
 restrict outgoing
                        Disallows remote remote queries to another host
 restrict incoming Disallows remote remote queries from other hosts
 DATE MODIFIED: 9.21.2008
restrict outgoing
restrict_ingoing
```

Note

The TM_config.txt configuration file is strictly optional. It can contain the restrict incoming or the restrict outgoing options.

You can secure the tm_config.txt by using the native Operating System file permissions (for example, chmod on UNIX) to make it read-only to all users except the administrator.

Appendix E – Message Triggering

Prior to Image Services Release 4.0, whenever the Image Services software detected a situation requiring some level of human intervention, the software generated an RSVP message that displayed in the Storage Library Control window. The System Administrator was then required to monitor the Storage Library Control window and respond to each message. In addition to RSVP messages, informative messages which did not require a user response were displayed in the INFO messages area of the Storage Library Control window.

Message triggering improves administrative responsiveness for RSVP and INFO messages. Whenever the system generates an RSVP or INFO message, it will continue to be displayed in the Storage Library Control window. However, with message triggering, you can also invoke a script or user-written program that can take specific actions in response to the RSVP message. For example, you can write a script that will notify the administrator or operator through e-mail that an RSVP/INFO event has occurred.

Note

You are responsible for the content and operation of your customized scripts.

Developing a Customized Script

You can have only one script per Storage Server and it will be launched for both RSVP and INFO message types. A script can be a normal shell script for UNIX or a batch command file for Windows. A program could be a C, C++, Java™ or any supported executable program type.

IS will invoke scripts in an asynchronous fashion and relevant information will be passed to the script or executable program as arguments.

Invoking the Script

The user interface for launching the script is:

<Script name><argument list>

Script Arguments

The following table describes the fixed list of arguments passed to the user-provided script. These arguments apply to both RSVP and INFO messages and can be distinguished by the type argument (under the field column). All arguments are inherently of string type but have a logical type as described in the type column below.

| Argument | Field | Туре | Valid Range | Comments |
|----------|---|----------------------|----------------------|---|
| 1 | type | integer | 1 - 2 | 1 = RSVP |
| | | | | 2 = INFO |
| 2 | error/event
indicator in
decimal format | integer | > 0 | Translated from the FileNet error/info tuple in decimal format. See appendix A |
| 3 | error/info indicator in string format | Variable string [42] | | Error tuple in <nn,nn,nn>. This three part error identifier can be used as input parameters to the fn_msg utility program to extract the specific error text descriptor.</nn,nn,nn> |
| 4 | ssn | integer | 1000 -
2147483646 | Unique System Serial Number |
| 5 | domain | Variable string [41] | | This includes the Domain:Organization name. |

| Argument | Field | Туре | Valid Range | Comments |
|----------|------------|----------------------|-------------------------|--|
| 6 | host_name | Variable string [50] | | This is the hostname or server name. |
| 7 | ip_address | Variable string [30] | | This is the IP address of the host. IP address format is: n.n.n.n |
| 8 | surface_id | integer | 3000 -
4,294,967,290 | -1 if not in use |
| 9 | media_type | integer | 1 - 30 | -1 if not in use See "Media Types" on page 635. Note that new media types may be supported in the future. Note that MSAR surfaces that have been converted from optical surfaces via the st_msar_convert background job will retain the optical media type. |

| Argument | Field | Туре | Valid Range | Comments |
|----------|--------------|---------------------|--|---|
| 10 | library_type | ibrary_type integer | 0 - 8 | -1 if not in use |
| | | | | 0 = Standard_osar |
| | | | | 1 = mini_osar |
| | | | | 2 = access_osar |
| | | | | 3 = 4500rapidchanger |
| | | | | 4 = HP osar |
| | | | | 5 = IBM osar |
| | | | | 6 = 6600rapidchanger |
| | | | | 7 = 8600rapidchanger |
| | | | | 8 = MSAR |
| | | | | Note that new library types may be supported in the future. |
| | | | | Note also that drive only devices (ODU) will have a slot number value of zero. See argument 15 below. |
| 11 | msar_path | Variable | | -1 = if it's not in use |
| | | string[256] | | Valid for MSAR file only. |
| | | | Contain MSAR surface file full pathname. | |
| | | | | For example: |
| | | | | On Windows, local drive, path-
name =
<drive>:\<path>\003000.dat.</path></drive> |
| | | | | |

| Argument | Field | Туре | Valid Range | Comments |
|----------|--------------|----------------------------|-------------|--|
| | | | | On Windows, remote shared drive, pathname is UNC name, pathname = \\ <path>\003000.dat</path> |
| | | | | On Unix, pathname = / <path>/003000.dat</path> |
| 12 | eject_status | integer | 1 - 2 | -1 = if it's not in use |
| | | | | 1 = media is automatically ejected due to MSAR error. |
| | | | | 2 = media to be ejected due to Optical error. |
| 13 | Library_id | A single char- | 'A' - 'H' | -1 if it's not in use |
| | | acter null ter-
minated | | Note that in the future more then 8 libraries may be supported. |
| 14 | drive_number | integer | 0 - 12 | -1 if it's not in use |
| | | | | Note that in the future more then 12 drives may be supported. |
| 15 | slot_number | integer | 1 - 2048 | -1 if it's not in use |
| | | | | Note that in the future more then 2048 slots may be supported. |
| | | | | Note also that drive only devices (ODU) will have a slot number of zero. This is how an ODU may be identified. |

| Argument | Field | Туре | Valid Range | Comments |
|----------|-----------------|--------------------------|-------------|---|
| 16 | RS232_fault_num | integer | | -1 if it's not in use |
| | ber | | | Valid for certain error for a serial interface (not SCSI) optical library. |
| 17 | msg_text | Variable
string[1024] | | RSVP or INFO text message describing the particular issue. These messages will be localized. Note that if there is no message in the ERM file, the "NO MESSAGE TEXT" string will be returned in this argument. This would indicate that the ERM file is not up to date. |

Error Tuples

The event indicator, argument 2 and argument 3 in the table above, is a FileNet error tuple. The *Image Services System Messages Hand*book lists some of the possible values that may be returned. Occasionally the event indicated can be serious, such as the file cannot be accessed or the Storage Library is broken. Other times, the error tuple is just an informational tuple, such as identifying of all media has been started. In addition to the event indicator, and to further distinguish an event, a library type field, argument 10, indicates whether an event is associated with an optical library or an MSAR library and the library number field indicates which library is having the problem. For example, if the error tuple states that a drive is down, the drive number field will reveal which drive and the library ID field will reveal which library malfunctioned.

New error tuples have been added to handle insertion RSVPs as well as a number of INFO messages.

Surface Ejection

The eject status flag, argument 12, with a value of 1 would mean that the surface has been automatically ejected.

Note

Only MSAR surfaces may be automatically ejected. This may happen if an inconsistent surface file or an older version of the surface file has been detected.

For an optical disk, the eject status may be used to indicate that a surface is dirty and needs to be cleaned. In this case, ejection always requires human intervention and the eject status flag is set to 2. When the user replies to the RSVP through the Storage Library Window, the optical surface will be ejected.

Message Text

The message text (argument 17) is a text description of the error tuple, which will look exactly the same as the text that is displayed when running

fn_msg <nn,nn,nn> .

These messages will be localized and will look similar to the RSVP or INFO text message generated by the Storage Library Control Program. The difference is that, the message generated in the Storage Library Control program may have additional information such as the library number, drive number, MSAR surface file name or surface id included in the message.

The following is a comparison of what is generated in the RSVP Storage Library Control Window and what is provided to the script program via the run-string arguments. This is only an example, but it

shows that all the same information is available to the script via the run-string arguments. Note also that msg_text argument is localized since it is extracted from the localized error file.

| Message in Storage Library Control Window | Information passed to the user-written script or program via the run-string parameters |
|---|--|
| Please load blank 1.3 GB Erasable media as sur- | -type = "1" <rsvp type=""></rsvp> |
| face 3000. Then respond to this message | -msg_text = "Insert a blank optical media to library" |
| | -library_type = "4" <hp library=""></hp> |
| | -library_id = "A" |
| | -media_type = "9" < 1.3GB Erasable> |
| | -surface_id = "3000" |
| | -error_tuple_text = "130,32,1" |
| MSAR surface 4000 file '/surf/004000.dat' has a | -type = "1" <rsvp type=""></rsvp> |
| permission problem. Change the permission before responding to this RSVP. | -msg_text ="Insufficient permission to open/read/write the file" |
| | -library_type = "8" < MSAR library> |
| | -library_id = "B" |
| | -surface_id = "4000" |
| | -msar_path = "/surf/004000.dat" |
| | -error_tuple_text= "202,100,11" |

| Message in Storage Library Control Window | Information passed to the user-written script or program via the run-string parameters |
|--|---|
| MSAR surface 5000 file '/surf/005000.dat' has | -type = "2" <info type=""></info> |
| been ejected from library c. It is out-of-sync with the database. Error <30,0,134> | -msg_text =" MSAR surface is out of sync with the database. Must restore the correct version of file before continuing, otherwise, some documents will exist in database but missing in the surface file." |
| | -library_type = "8" < MSAR library> |
| | -library_id = "C" |
| | -surface_id = "5000" |
| | -msar_path = "/surf/005000.dat" |
| | -ejected = "1" < file is automatically ejected> |
| | -error_tuple_text = "30,0,134" |
| Library d, Drive 1: Surface 6000 Error <30,0,117> | -type = "1" <rsvp type=""></rsvp> |
| WPC error. Please clean media and retry writing to same media. | -msg_text = "WPC error. Please clean media and retry writing to same media.If WPC errors continue, disable this media so software writes to a different media. If it fails again, take the drive offline for repair and then manually perform a FORCE WPC operation via the drive's panel." |
| | -library_type = "0" < FileNet library> |
| | -library_id = "D" |
| | -surface_id = "6000" |
| | -ejected = "2" <surface be="" ejected="" needed="" to=""></surface> |
| | -error_tuple_text = "30,0,117" |

Media Types

The following table describes the currently supported media types.

| Media Type | Media Type Description | | |
|------------|-----------------------------|--|--|
| 0 | Unknown media type | | |
| 1 | Hitachi 12" 2.6 GB WORM | | |
| 2 | Hitachi 5" 600 MB WORM | | |
| 3 | Hitachi 5" 600 MB Erasable | | |
| 4 | Philips 12" 2 GB WORM | | |
| 5 | Hitachi 12" 7 GB WORM | | |
| 6 | Philips 12" 5.5. GB WORM | | |
| 7 | HP 5" 650 MB Erasable | | |
| 8 | HP 5" 650 MB WORM | | |
| 9 | Standard 5" 1.3 GB Erasable | | |
| 10 | Standard 5" 1.3 GB WORM | | |
| 11 | Philips 12" 12 GB WORM | | |
| 12 | Standard 5" 2.6 GB Erasable | | |
| 13 | Standard 5" 2.6 GB WORM | | |
| 14 | IBM 5" 2.6 GB Ablative WORM | | |
| 15 | Standard 5" 5.2GB Erasable | | |
| 16 | Standard 5" 5.2GB WORM | | |
| 17 | IBM 5" 5.2GB Ablative WORM | | |
| 18 | Philips 12" 30 GB WORM | | |
| 19 | Standard 5" 9.1 GB Erasable | | |
| 20 | Standard 5" 9.1 GB WORM | | |
| 21 | MSAR 1GB | | |
| 22 | MSAR 2GB | | |
| 23 | MSAR 4GB | | |

| Media Type | Media Type Description | |
|------------|------------------------|--|
| 24 | MSAR 8GB | |
| 25 | MSAR 16GB | |
| 26 | MSAR 32GB | |
| 27 | 5" UDO 30 GB Erasable | |
| 28 | 5" UDO 30 GB WORM | |
| 29 | 5" UDO 60 GB Erasable | |
| 30 | 5" UDO 60 GB WORM | |

Script Example

The following is a sample UNIX C shell script that provides an e-mail notification to a system administrator when an RSVP or INFO event has occurred.



For UNIX platforms, this script is available in /fnsw/etc/sample/rsvp_csh.



For Windows platforms a similar script that only displays all arguments (no e-mail notification) is available in /fnsw/etc/sample/rsvp.bat.

^{#!/}bin/csh -f

[#] RCI 2810 - RSVP/INFO trigger feature

[#] This script is launched whenever an RSVP or INFO message is generated

[#] by the IS software and the RSVP/INFO triggering feature is turned on.

[#] After it is invoked, this script will first print all the arguments to a

[#] temporary file. Then it invokes the 'e-mail' program to send that file,

[#] and then to remove the temporary file.

[#] To send the e-mail to your e-mail address, simply change the e-mail

[#] address in the mail command and uncomment that line of code by removing

[#] the pound sign (#). To remove the temporary file, simply uncomment the

```
# last line of code by removing the pound sign.
# The temporary file is defined in the variable 'temp out', which is in
# /fnsw/local/logs/rsvp.out$$. The directory /fnsw/local/logs/rsvp
# must be created by the user before executing this script, and
# the script must have read and write permission to the directory.
# To change the 'temp out' to a different file system, modify your
# temporary file name path.
# Add execution permission to this script (i.e. chmod ugo+rx rsvp csh)
set type
                 = $1
set err dec
                = $2
set err
                = $3
set ssn
                = $4
set domainorg
                = $5
set host name
               = $6
set ip addr
                = $7
set surface_id = $8
set media type = $9
set library type = $10
set msar file name = $11
set eject_status = $12
set library
                = $13
set drive_num
                = $14
set slot_num = $15
set fault num
                = $16
set msg_txt = "$17"
set opt disk type = 0
set temp_out = /home/fnsw/rsvp/rsvp.out$$
echo "DATE TIME STAMP = `date`" >> $temp out
                   = `hostname`" >> $temp out
echo "HOSTNAME
echo >> $temp out
if (type == 1) then
  echo "RSVP MESSAGE" >> $temp_out
```

```
else
   echo "INFO MESSAGE" >> $temp out
endif
echo "=======" >> $temp out
echo "ERROR/INFO TUPLE = $err_dec (decimal)" >> $temp_out
echo "ERROR/INFO TUPLE = $err" >> $temp out
echo "SSN = $ssn" >> $temp_out
echo "DOMAIN_ORG = $domainorg" >> $temp_out
echo "HOST_NAME = $host_name" >> $temp_out
echo "IP ADDRESS = $ip addr " >> $temp out
if ($surface id != -1) then
   echo "SURFACE ID = $surface id" >> $temp out
endif
switch ($media type)
   case -1:
           breaksw
   case 1 :
           echo 'MEDIA TYPE = Maxell 12" 2.6 GB WORM' \
            >> $temp out
            set opt disk type = 1
            breaksw
   case 2 :
           echo 'MEDIA TYPE = Maxell 5" 600 MB WORM' \
            >> $temp out
            set opt disk type = 1
            breaksw
   case 3 :
           echo 'MEDIA_TYPE = Maxell 5" 600 MB Erasable' \
            >> $temp out
            set opt disk type = 1
            breaksw
   case 4 :
           echo 'MEDIA TYPE = Plasmon/Philips 12" 2 GB WORM' \
```

```
>> $temp_out
         set opt disk_type = 1
         breaksw
case 5 :
        echo 'MEDIA TYPE = Maxell 12" 7GB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 6 :
        echo 'MEDIA TYPE = Plasmon/Philips 12" 5.5 GB WORM' \
         >> $temp_out
         set opt disk type = 1
         breaksw
case 7 :
        echo 'MEDIA TYPE = HP 5" 650 MB Erasable' \
         >> $temp_out
         set opt disk type = 1
         breaksw
case 8 :
        echo 'MEDIA TYPE = HP 5" 650 MB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 9 :
        echo 'MEDIA TYPE = Standard 5" 1.3 GB Erasable' \
         >> $temp_out
         set opt disk type = 1
         breaksw
case 10:
        echo 'MEDIA_TYPE = Standard 5" 1.3 GB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 11:
        echo 'MEDIA TYPE = Plasmon/Philips 12" 12GB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
```

```
case 12:
        echo 'MEDIA TYPE = Standard 5" 2.6 GB Erasable' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 13:
        echo 'MEDIA_TYPE = Standard 5" 2.6 GB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 14:
        echo 'MEDIA TYPE = IBM 5" 2.6 GB Ablative WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 15:
        echo 'MEDIA TYPE = Standard 5" 5.2 GB Erasable' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 16:
        echo 'MEDIA TYPE = Standard 5" 5.2 GB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 17:
        echo 'MEDIA TYPE = IBM 5" 5.2 GB Ablative WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 18:
        echo 'MEDIA TYPE = Plasmon/Philips 12" 30GB WORM' \
         >> $temp out
         set opt disk type = 1
         breaksw
case 19:
        echo 'MEDIA TYPE = Standard 5" 9.1 GB Erasable' \
```

```
>> $temp out
            set opt disk type = 1
            breaksw
  case 20:
           echo 'MEDIA TYPE = Standard 5" 9.1 GB WORM' \
            >> $temp out
            set opt disk type = 1
            breaksw
  case 21:
           echo 'MEDIA TYPE = FileNet MSAR 1 GB' \
            >> $temp out
            breaksw
  case 22:
           echo 'MEDIA TYPE = FileNet MSAR 2 GB' \
            >> $temp_out
            breaksw
  case 23:
           echo 'MEDIA TYPE = FileNet MSAR 4 GB' \
            >> $temp out
            breaksw
  case 24:
           echo 'MEDIA TYPE = FileNet MSAR 8 GB' \
            >> $temp_out
            breaksw
  case 25:
           echo 'MEDIA_TYPE = FileNet MSAR 16 GB' \
            >> $temp out
            breaksw
  case 26:
           echo 'MEDIA TYPE = FileNet MSAR 32 GB' \
            >> $temp_out
            breaksw
  default
           echo 'MEDIA TYPE = Unknown media type' \
            >> $temp_out
            breaksw
endsw
```

```
if ($surface id != -1) then
  if ($library type == 8) then
     if ($opt disk type == 1) then
        echo '
                                  (This is an optically converted MSAR
surface) ' \
            >> $temp out
     else
        echo '
                                  (This is an MSAR surface)' \
            >> $temp out
     endif
     echo "FILE NAME = $msar file name" >> $temp out
  else
     echo '
                              (This is an optical surface) ' \
            >> $temp out
      if (fault num != -1) then
         echo "RSS232 FAULT NUM = $fault num" >> $temp out
     endif
  endif
endif
switch ($library_type)
    case 0:
            echo 'LIBRARY TYPE = Standard osar' \
             >> $temp out
             breaksw
    case 1:
            echo 'LIBRARY TYPE = mini osar' \
             >> $temp out
             breaksw
    case 2:
            echo 'LIBRARY TYPE = access osar' \
             >> $temp out
             breaksw
    case 3:
            echo 'LIBRARY TYPE = 4500 Rapidchanger' \
             >> $temp out
             breaksw
```

```
case 4:
            echo 'LIBRARY TYPE = HP osar' \
             >> $temp out
             breaksw
    case 5:
            echo 'LIBRARY TYPE = IBM osar' \
             >> $temp out
             breaksw
    case 6:
            echo 'LIBRARY TYPE = 6600 Rapidchanger' \
             >> $temp out
             breaksw
    case 7:
            echo 'LIBRARY TYPE = 8600 Rapidchanger' \
             >> $temp_out
             breaksw
    case 8:
            echo 'LIBRARY TYPE = MSAR' \
             >> $temp out
             breaksw
    default
            breaksw
endsw
if (\$slot num == 0) then
  echo 'The above library is an Optical Disk Unit. (No gripper)' \
   >> $temp out
endif
if ($library_num != -1) then
  echo "LIBRARY_NUMBER = $library_num" >> $temp_out
endif
if (\$drive num != -1) then
  echo "DRIVE NUMBER = $drive num" >> $temp out
endif
```

```
if (\$slot num != -1) then
   echo "SLOT NUMBER
                          = $drive num" >> $temp out
endif
echo >> $temp out
if ($eject status == 1) then
   echo 'MSAR surface is automatically ejected due to error' \
        >> $temp out
else if ($eject status == 2) then
   echo 'Optical surface needed to be ejected, please respond to RSVP' \
        >> $temp out
endif
echo >> $temp out
echo MESSAGE
                      = $msq txt >> $temp out
echo >> $temp out
#mail SysAdmin@yourcommany.com < $temp out</pre>
#rm -f $temp out
```

Based on the above script, if there is a disk insertion RSVP posted on the Storage Library Control window, the system administrator will receive an e-mail similar to the following:

DATE_TIME_STAMP = Wed Feb 5 19:21:53 PST 2003

HOSTNAME = liberty

INFO MESSAGE =========

ERROR/INFO TUPLE = 2183135233 (decimal)

ERROR/INFO TUPLE = 130,32,1

= 3287 SSN

DOMAIN_ORG = vicenza:FileNet HOST_NAME = liberty

IP_ADDRESS = 10.14.102.8
SURFACE ID = 3768

SURFACE ID = 3768

MEDIA_TYPE = Standard 5" 5.2 GB Erasable

(This is an optical surface)

LIBRARY TYPE = HP osar

LIBRARY NUMBER = B

MESSAGE = Insert a blank optical media to library

If there is a message in the Storage Library Control window indicating an MSAR surface not found, the system administrator will receive an email similar to the following:

DATE_TIME_STAMP = Thu Feb 6 12:21:51 PST 2003

HOSTNAME = vicenza

INFO MESSAGE

=========

ERROR/INFO TUPLE = 3395551242 (decimal)

ERROR/INFO TUPLE = 202,100,10

= 3287 SSN

= vicenza:FileNet DOMAIN ORG

HOST NAME = vicenza

= 10.14.102.102 IP ADDRESS

= 3000SURFACE ID

MEDIA_TYPE = FileNet MSAR 1 GB

(This is an MSAR surface)

LIBRARY TYPE = MSAR LIBRARY_NUMBER = A

MSAR surface is automatically ejected due to error

MESSAGE = FCL: specified file, device or path does not exist

Appendix F – Multicultural Support

Getting started

The guidelines in this appendix provide information for system administrators to use when installing and configuring IBM FileNet Image Services (IS) in non-English environments.

Language support for Image Services is as follows:

Table 1: Image Services Language Support

| Activity | Supported Languages | |
|----------------|--|--|
| Installation | English | |
| Configuration | English, French, German | |
| Error messages | English, French, German | |
| Indexing | Multiple languages as described in
"Character set support" on
page 655 | |

Operating systems

Servers that you purchase are preinstalled with the language that you specified at the time of order. In certain cases, it is possible to support several languages in addition to the pre-installed version. For example, if you have a server that uses the Windows operating system installed in English, you can install support for both the Chinese and Greek character sets.

Important

All Image Services and Image Services Toolkit (ISTK) servers and clients must be configured to use compatible character sets at both the operating system level and the relational database level.

The Image Services software interface is available in English, French and German.

- For Windows, you can use a number of languages for the localized operating system. See your Microsoft Windows documentation for more information.
- For UNIX, use the values in the following table to display the appropriate interface.

Table 2: Language variables for French and German

| Language | AIX | HP-UX | Solaris |
|----------|-------|----------------|---------|
| French | fr_FR | fr_FR.is088591 | Fr |
| German | de_DE | de_DE.iso88591 | De |

Setting the Language variable

The LANG variable specifies the language, or locale, that the software uses to display messages and instructions, collating sequences, date formats, and so on.

Note

The Image Services software always defaults to English if the LANG environment variable is not set to any of the supported languages: English, French, or German.

Different operating systems use different locale names to specify the same language. Skip to the section for your server:

- "Setting the LANG variable on AIX" on page 649
- "Setting the LANG variable on HP-UX" on page 651
- "Setting the LANG variable on Solaris" on page 652
- "Setting the LANG variable on Windows Server" on page 653

See the documentation for your operating system for additional information on managing the LANG environment variable.

Setting the LANG variable on AIX

Set up the LANG variable in your shell environment files using your preferred text editor (such as vi). From your HOME director, edit the necessary files by completely the steps in one of the following:

Important

Verify that the language you select here matches the language that was selected when the relational database software was installed.

For **sh** and **ksh**, add the following lines to your .profile file:

LANG=<language> export LANG

where <language> is the native language (locale) in which you want your server to operate. For example:

> French is fr FR German is de DE

For French, you would add these lines:

LANG=fr FR export LANG

Or for German you would add these lines:

LANG=de DE export LANG

For **csh**, add the following line to your .login file:

setenv LANG < language>

where <language> is the native language (locale) in which you want your server to operate.

For French, you would add this line:

setenv LANG fr FR

Or for German you would add this line:

setenv LANG de DE

Note

When you install or upgrade Image Services software on an AIX server that is set to any Japanese locale such as ja_JP, you must set the LANG environment variable to en_US (English) before running the installation program. Setting the LANG variable to en US prevents garbled characters from displaying during the installation. After the IS installation is finished, you can reset the LANG variable to its original locale, such as ja JP.

Setting the LANG variable on HP-UX

Set up the LANG variable in your shell environment files using your preferred editor (such as vi). From your HOME directory, edit the necessary files by completing one of the following steps, depending upon the shell you are using:

For **sh** and **ksh**, add the following lines to both your .profile and .vueprofile files:

```
LANG=<locale>.iso88591
export LANG=C
```

where <locale> is the abbreviated language name your Xstation is dedicated to. For example:

```
French is fr FR.iso88591
German is de DE.iso88591
```

For French, you would add these lines:

```
LANG=fr_FR.iso88591
export LANG=C
```

Or for German you would add these lines:

```
LANG=de DE.iso88591
export LANG=C
```

For **csh**, add the following line to both your .login and .vueprofile files:

```
seteny LANG < locale > . iso88591
```

where <locale> is the language your Xstation is dedicated to.

For French, you would add this line:

setenv LANG fr FR.iso88591

Or for German you would add this line:

setenv LANG de DE.iso88591

Setting the LANG variable on Solaris

Set up the LANG variable in your shell environment files using your preferred editor (such as vi). From your HOME directory, edit the necessary files by completing one of the following:

Important

Verify that the language you select here matches the language that was selected when the relational database software was installed.

For **sh** and **ksh**, add the following lines to your .profile file:

LANG=<language> export LANG

where <language> is the native language (locale) in which you want your server to operate. For example:

> French is **Fr** German is **De**

For **csh**, add the following line to your .login file:

setenv LANG < language>

where <language> is the native language (locale) in which you want your server to operate.

Non-localized environment

To use 8-bit characters, the Solaris 10 operating system software must run in a European locale. However, European locales require installation of localization feature packages.

If these localization packages are not available to you, use the following work-around:

For sh and ksh, enter:

LC CTYPE=iso 8859 1 export LC_CTYPE

For **csh**. enter:

setenv LC CTYPE iso 8859 1

Note

The ISO 8859 1 setting is not valid for the OpenWindows locale X resources. The ISO 8859 1 setting is not a full locale setting, but a setting for the LC_CTYPE category only. DO NOT set the LANG environment variable to ISO 8859 1.

Setting the LANG variable on Windows Server

On Windows Server systems, the Image Services applications determines the desired language based on the following settings:

- The LANG environment variable
- The default OS locale (See your OS documentation for details.)

The LANG variable is an optional setting, but if you set it, it will always override the default OS locale. You can set it using the System Properties dialog, or at a command prompt in a CMD window.

For example, launch a CMD window and enter:

setenv LANG < language>

where <language> is the native language (locale) in which you want your server to operate. For example:

> French is fra German is deu

Important

Verify that the language you select here matches the language that was selected when the relational database software was installed.

Relational databases

When the relational database management system (RDBMS) is installed, the language and character set for the database defaults to the language you selected for the operating system.

CAUTION

To change the database character set after the RDBMS is installed, refer to the appropriate RDBMS documentation.

Note also that the NLS LANG environment variable determines if the database will perform any character translations before storing characters in the database or after retrieving characters from the database. For optimum database performance, NLS_LANG is set to the same characters as the "Database Character Set."

Image Services configuration

In the IS System Configuration Editor (fn_edit), select the System Attributes tab and set the appropriate default character set and former character set. If Image Services needs to be configured to process Japanese characters, select either Shift Japanese or Japanese EUC for both the default and former character sets, depending on your usage. The IS configuration editor provides both single-byte and double-byte (or multi-byte) character set support. The character set must be selected before initializing the IS databases.

Character set support

IS supports single-byte character sets (SBCS) and multi-byte character sets (MBCS). However, some limitations exist for both homogeneous and heterogeneous configurations.

The default IS character set is configured through fn edit by specifying an International Organization for Standardization (ISO) character encoding. Regardless of whether the IS platform is UNIX-based or Windows-based, the default character set is always defined as an ISO character encoding. Therefore, IS assumes that all data sent to it by clients is in ISO format, rather than in a Windows code page (CP) format.

Single-byte character sets

SBCSs are based on 8-bits and can represent up to 256 characters. These characters are often referred to as a code page on Windows servers. For example, the default code page on Windows servers with IS is CP1252 (also called the ANSI code page). The equivalent ISO character encoding on UNIX servers is ISO 8859-1.

The following table lists all the supported SBCS Windows Code Pages and their equivalent ISO encoding on Unix operating systems.

Table 3: Single-Byte Character Sets

| Windows Code
Page | fn_edit Name | Language Group | NLT Map
Provided |
|----------------------|-----------------------|-----------------------------|---------------------|
| 1252 | ISO 8859-1 | Latin 1
(Western Europe) | Yes |
| 1250 | ISO 8859-2 | Latin 2
(Central Europe) | Yes |
| 1257 | ISO 8859-4 | Scandinavia and Baltic Rim | Yes |
| 1251 | Cyrillic (ISO 8859-5) | Cyrillic (Slavic) | Yes |
| 1256 | Arabic (ISO 8859-6) | Arabic | Yes |
| 1253 | Greek (ISO 8859-7) | Greek | Yes |
| 1255 | Hebrew (ISO 8859-8) | Hebrew | Yes |
| 1254 | Turkish (ISO 8859-9) | Turkish | Yes |
| 874 | Thai TIS 620-2533 | Thai | No |

Note

The IS National Language Translation (NLT) subsystem automatically translates between single-byte Windows and UNIX code pages when Windows IS clients communicate across the network. The Thai character set is not fully supported in heterogeneous environments because a translation map is not available.

Double-byte and multi-byte character sets

In addition to single-byte character sets, IS supports several doublebyte character sets (DBCS) and multi-byte character sets (MBCS). The terms DBCS and MBCS are often used interchangeably, but MBCS is a more generalized form of DBCS. In a multi-byte character set, the

number of bytes per character can vary from 1 to 3. In both types of character set, each character can be represented by either one or two bytes. MS932 is an example of a DBCS, while eucJP is an example of a MBCS.

Important

The IS and RDBMS software has specific maximum length requirements for many of its parameters and database fields. When the IS system is configured with a DBCS or MBCS, the maximum number of characters is correspondingly fewer.

For example, if a database field has a maximum length of 24 (singlebyte) characters, the field can hold only 12 double-byte characters, and only 8 three-byte characters.

IS does not natively support any Wide Character Sets (WCS), which use a fixed number of bytes for each character in the set. For example, two, three, or four bytes per character. If a custom application manipulates strings internally using Unicode, the application is responsible for translating those strings to either a single-byte or double-byte character set before passing them to an IS application.

You can specify the following DBCS and MBCS through fn_edit:

Table 4: Double-Byte and Multi-Byte Character Sets

| Windows Code
Page | fn_edit Name | Language | NLT Map
Provided |
|----------------------|---------------------------------------|----------|---------------------|
| 932 | Shift Japanese Industrial
Standard | Japanese | No |
| 936 | Chinese Simplified EUC (China) | Chinese | No |
| 949 | Korean 5601 | Korean | No |

| Windows Code
Page | fn_edit Name | Language | NLT Map
Provided |
|----------------------|----------------------------------|----------|---------------------|
| 950 | Chinese Traditional EUC (Taiwan) | Chinese | No |
| | Japanese Extended UNIX
Code | Japanese | No |

Table 4: Double-Byte and Multi-Byte Character Sets, Continued

- Japanese EUC
- Chinese Simplified EUC (China)
- Chinese Traditional EUC (Taiwan)
- Korean 5601

Note

When the IS default character set is a DBCS or MBCS, the NLT subsystem does not perform any character translation across the network between IS clients and servers.

Heterogeneous environments

For homogeneous IS environments, such as Windows to Windows or UNIX to UNIX, there are no known issues with supporting DBCS or MBCS character sets because no character translation occurs between clients and servers.

For heterogeneous IS environments, translation between equivalent character sets is limited to those defined in "Table 3: Single-Byte Character Sets" on page 656, and for which an NLT map is available.

Known issues and limitations

Image Services does not support GB 18030 (Chinese Standard) characters that are four bytes long.

In heterogeneous environments, such as systems that include both Windows and UNIX servers, correct translation is not guaranteed. For example, in the case of a Japanese Windows system, the 932 character set is based on Shift JIS. Solaris supports EUC for Japanese, so character translation between a Japanese Windows server and a Japanese Solaris server might not be accurate.

However, the AIX operating system supports the Shift JIS standard and is therefore more likely to provide correct translation with a Windows platform, which also supports Shift_JIS.

To minimize compatibility issues in heterogeneous environments, all IS and ISTK clients must be configured to use compatible character sets at the operating system and database levels.

Appendix G – Some Practical Limitations of Image Services Components

Since its inception, Image Services software has been a top performer. Its basic design and strength is focused on committal and retrieval of document images or similar media. As new media types emerge, such as video streams, DVD data and blob data types, customers create new ways of using IS software. One of these new uses involves using IS as a mass data storage engine. This appendix describes limitations and the accompanying impact on performance that clients should be aware of when using IS as a mass storage engine.

Overview

When building an end-user application on top of Image Services and its ISTK(WAL) and IDM APIs, there are a number of considerations involved in assessing the maximum size a document can have. It is important to consider the practical limits of all the components involved in the chain to both ingest and retrieve documents, and take the most limiting factor in the chain as the upper limit.

Typically the following components will be included in the chain:

- Client Application
- IDM API
- ISTK (WAL) API
- Courier (Network) Services Layer

Image Services Internal Constructs (page, document, pages per document, all data structures)

End-user client applications have been the most limiting factor due mainly to the inherent limitations of both the client machine hosting the application and the application itself. Typical applications may be viewers or document editors. When handling images above a a few dozen MB in size, the end-user experience slows down to the point where it can become impractical.

In response to these impractical end-user client applications, new cases involving engines that process large document objects are emerging. These newer engines need to take into consideration that practical limits of Image Services server components. The following guidelines should be observed.

Document size

The practical size of a document should not exceed 30 MB per balanced page. All IS document data should be stored on evenly divided pages if possible. Data distributed evenly across multiple pages has advantages. For example, comparing a document of 1 page (500 mb) with a document of 500 pages (1 mb each), a read of a full page would take 1 second as opposed to 500 seconds. 500 seconds is the approximate time required to transfer 500 mb of data on a standard ethernet network (1 mb per second). However, if the client application needs all the 500 mb first, then it would still require 500 seconds to read in all 500 pages. Designing your document data storage using evenly distributed pages serves as a 'fetch on demand' method. Storing document objects using fetch on demand makes your application more scalable and robust. The maximum number of pages allowed per document is 1000. The maximum size of a page should not greater than

30 mb. 1000 x 30 gives you a capable document size of 30 gb. If this size is still not large enough, please contact your service representative.

Storage Device Capacity

Verify that your storage device capacity is large enough to store large document objects.

For MSAR clients, verify that your free hard space is large enough to handle large document objects. For OSAR clients, please be aware that any single document must be contained in a single surface. A single document cannot be spanned across multiple surfaces.

Network Speed Capacity

All end-user client applications are ISTK(WAL) based. Therefore, these applications use Courier Services Layer and TCP/IP for all Ethernet network data transfers. it is critical that all applications should retrieve documents with low rates of data transfer, if possible. The performance of your client application is critically linked to the amount of data transferred through your network. On a standard Ethernet network, the transfer rate is roughly 1 MB per second.

Glossary

This Glossary identifies other glossary entries in italics.

administrative domain

Any user with administrative attributes can administer his own administrative group and any other administrative group he is a member of. This is the user's administrative domain.

administrative user

An administrator is a *user* assigned special privileges to perform security tasks affecting users, groups, and devices. The four administrative attributes are: supervisor, principal, group, and password. By assigning different combinations of administrative attributes to different users, you can provide checks and balances in your security system.

ageable cache

Ageable cache is time-limited storage on magnetic media. Objects remaining in ageable cache past a specified time are eligible for deletion if space is needed to store other objects. See page cache.

annotations

Annotations, created in the FileNet Display application, function like stick-on notes for your documents. See your desktop application's user's guide for details on creating and displaying annotations.

API

An API (Application Programming Interface) is a generic term for any language and format used by one program to help it communicate with another program. In imaging, a vendor can provide an API that enables programmers to repackage or recombine parts of the vendor's imaging system.

archive

Archiving is the process of moving obsolete or unused documents out of the index database and to an archive database. Image Services for AIX offers archive tools.

Application Executive

Application Executive is a program you use to access the FileNet applications available for system management. Before starting Application Executive, you must start the FileNet *Task Manager*.

application server

An application server provides extra power to certain applications. For example, Document Entry might need an extra server to process batches or manage a large number of WorkFlo gueues.

Background Job Control

Background Job Control is the FileNet application you use to start, control, and monitor background jobs.

background jobs

Background jobs are time-intensive FileNet software functions designed to run in the background, including copying *media*, copying annotations from magnetic disk to storage media, importing a specified set of documents from storage media, consolidating media, and finding open documents on storage media. While these jobs are processing, you can perform other administrative tasks.

backup

A backup is a duplicate copy of data placed in a separate location from the main body of data (for example, on a tape, on a disk in a vault) to guard against total loss in the event the original data somehow becomes inaccessible.

BES

BES (Batch Entry Server/Services) is an application server on a large FileNet system that provides basic services, as well as Object Entry, Cache, and OSAR Management services.

B-tree

A B-tree is a database structure using a binary tree. One tree is built for each retrieval key, maintaining ordered pointers to documents containing the key to speed retrievals.

batch

A batch is a collection of *images*, all of the same *document class*, that are entered into the FileNet system as a group.

batch cache

This *cache* contains batches of *images* as they enter the system, typically through scanning or from the COLD application.

cache

Cache is the magnetic disk space used to store documents on the way to and from storage media (and can act as permanent storage when you do not use optical storage media). Portions of the cache storage are allocated to the different cache types (referred to as logical caches). See ageable cache, batch cache, folder notes cache, ICR cache, page cache, print cache, Revise cache.

Cache Export/Import

The Cache Export/Import program backs up your cache storage independently of system backups. This program enables you to export your cache data to import on another Image Services server.

See the System Administrator's Companion for UNIX or System Administrator's Companion for Windows Server for information about the Cache Export/Import Program. To download these documents from the IBM support page, see "Accessing IBM FileNet documentation" on page 35.

cataloging

Cataloging is the process of writing index information to the *index* database during committal.

client

A client is software running on a station in the system that invokes a service. The service runs on a server, which can be on another station or on the same station as the client.

clustering

Clustering directs the FileNet system to store all documents with a common index value (for example, the same loan number) in a reserved space on particular media.

COLD

The COLD (Computer Output to Laser Disk) application reads data from magnetic tapes, converts the data into FileNet documents, and automatically indexes and commits the documents (via fast batch committal).

committal

Committal involves three processes. Assembly is the stage at which images become documents. Cataloging is the process of adding index values to the index database. Migration is the process of writing a document's optical disk information to the Permanent database and of the document to optical or magnetic disk. Assembly and cataloging take place at the same time. Migration takes place when it is convenient for the system based on system priorities.

consolidate media

Consolidating media is the process of copying all undeleted documents to other media to make room in the storage library for current documents.

core files

Core files (known as Dr. Watson files on Windows Server systems and core files on UNIX platforms) are files that are created due to system or application program errors. Core files often contain useful information that can help your service representative solve system problems.

cross-system committal

Cross-system committal provides a method of committing documents to a different Image Services server than the server on which the batch entry services (BES) application resides.

data object

Data objects include documents, folders, and annotations. Security is assigned to these objects through the groups, which are granted read, write, and append/execute privileges.

database

A database is a collection of logically related records or files. The FileNet system uses two types of databases: a third-party RDBMS (relational database management system) for index data and MKF (multi-keyed file) databases for document addresses and work in progress.

Database Maintenance

Database Maintenance is the FileNet application you use to create and maintain indexes, document classes, and media families, and archived information. Database Maintenance also provides reports on indexes, document classes, and media families.

device

See security device.

dialog

See dialog box.

dialog box

A dialog box is a popup window containing fields that require input from the user.

distributed file system

A distributed file system allows files managed by the operating system to be distributed across multiple stations on the same local area network. The location of the files is transparent to application programs.

DMA

The DMA (Document Management Alliance) consists of multiple user and vendor companies working together as an AIIM task force to define the DMA specification. The DMA specification defines the search and retrieval software components of large, multi-vendor document management systems.

doctaba

Doctaba is the principle table in the Index DB which stores one data record for each document stored on the IS system.

document

Documents can be images, text, forms, mixed (combinations of types), imported DOS files stored on the FileNet system's storage media, or overlays checked in by Revise users. See *committal*.

document class

A document class describes the scanning, indexing, and security characteristics of a group of documents.

document header file

A document header file, written on storage media, contains a record of the documents committed over a period of time.

document ID

A document ID is a unique number the FileNet system assigns to a document. This number is used in the tables and databases in the system to keep track of documents.

document locator server

The document locator server is a *storage library server* that includes the complete permanent database that maps each document number into two media locations. Other storage library servers contain small permanent databases that store this information until enough accumulates to transfer to the document locator server. The document locator server is the only server running document services.

document services

Document services is the part of the FileNet software that stores and retrieves documents on the storage library server.

domain name

The domain name is the second part of the *NCH* resource name. The domain name is the system name, which is determined by you and set up by your service representative during FileNet system configuration.

Dr. Watson files

See core files.

drive

A drive is the physical hardware necessary for reading and writing media

dump files

See core files.

event logs

Event logs, created daily, contain error messages and entries for activities occurring in the system. Use the Task Manager to review event logs for FileNet software.

expiration

Expiration makes a user, group, or terminal unusable after a time set by an administrator has been exceeded.

extended membership

Through extended *membership*, a user inherits the *permissions* of all the extended groups and can perform tasks that require permissions beyond those explicitly assigned to his group.

fast batch committal

Fast batch committal is a quick way of moving documents from batch cache to page cache (making the documents available for retrieval). In a fast batch committal environment, all documents and all pages of the documents for the batch are in one cache object, so the entire batch is committed in a single operation.

field

In windows and dialog boxes, fields are unprotected areas where you can enter data.

folder

A folder is a logical grouping of document images. A folder has a specified set of retention, disposition, and filing parameters, and a name, pathname, and ID number.

folder notes cache

If your system includes FolderView, you can store folder notes in cache on the local PC, a network PC server, or a FileNet server. The official name of the cache on the Image Services server is folder_cache. In the System Monitor's Magnetic Disk Cache Info report, folder notes cache is referred to as Folder View Cache.

foreign media

Foreign *media* is media that was formatted and written on a system other than yours. When you incorporate media from another system, that media is listed as foreign in certain *SLC* reports.

form

In data and image processing systems, a form is a formatted display for data entry consisting of protected prompts and unprotected entry areas (fields).

function

As used by the Security Administration program, a function is an option on a menu that you can secure so that only certain users can perform it.

gripper

A gripper is the part of the robotic arm in a storage library that grasps the media.

group

A group is a security object to which one or more users, devices, or other groups can be assigned. You make data objects accessible by assigning the name of the group that can read, write, or append/ execute them.

GUI

GUI is an acronym for graphical user interface.

GUID

GUID is an acronym for Globally Unique IDentifier. GUIDs are DMAcompliant, 16-byte integers used to uniquely identify each element transported over a network. The system ensures unique GUID assignments by automatically generating this integer using an algorithm based on the system's network card MAC address and a format that complies with the specifications provided for the system's platform.

highlights

Highlights, created in the FileNet Display application, function like highlighting accent markers for your documents. See your desktop application's user's guide for details on highlights.

ICR cache

On systems that include WorkFlo/ICR, cache services creates and stores each batch of images as one cache object. The objects are stored in ICR (intelligent character recognition) cache until written to storage media.

Image Services

A set of servers and services providing a single document image database. The database includes a single index database, a single document locator database, and the collection of document images on storage media.

index

An index contains the information used for retrieving documents. All index information is stored in the index database and also on storage media in page zero of the document. Later, when you need to look at the document, the FileNet software looks in this database for index information that satisfies a retrieval query.

index database

The index database, an *RDBMS* database (Oracle, DB2 or Microsoft SQL Server), contains document and folder information and can contain WorkFlo queues.

informational index

Informational index is a term used to refer to indexes that are not set up as retrieval keys.

interleaving

Interleaving is the process of writing the A sides of several media before returning to write the B sides. Interleaving media keeps the most recently committed documents immediately available for retrieval without flipping the disk to write on the other side.

jukebox

A jukebox is a device, also know as a storage library, that holds multiple optical discs and one or more disc drives. A jukebox can swap discs in and out of the drive as needed.

library

See storage library.

library server

See storage library server.

locked objects

Images not yet written to storage *media* are "locked" in *cache* and cannot be deleted.

logical cache

See cache.

magnetic disk

Magnetic disk, usually an internal hard disk on your system, is where the Image Services software, cache, databases, etc. are stored.

magnetic disk cache

See cache.

margin notes

Margin notes, created in the FileNet Display application, are colorhighlighted notes you can place in your documents. See your desktop program's user guide for details on margin notes.

mask

A mask is a template that displays elements such as dates and numbers in a particular format. For example, mm/dd/yyyy is a mask that tells the system to display the date like this: 12/02/1999.

media

Media is any material on which data is stored (magnetic disk, optical disk, magnetic tape). We usually refer to optical disks as storage media.

media family

The media family defines what type of storage *media* the *document* class uses. In general, the media family controls which media surfaces will be used by the document classes that use the family.

membership

Users and devices can be assigned membership in one or more groups. A user or device inherits the permissions of all groups of which the user or device is a member. A group can be assigned membership in one or more groups. The group inherits the permissions of all groups of which it is a member.

migration

Migration is part of the Committal process. Migration writes documents from magnetic disk to storage media.

MKF

The MKF (Multi-Keyed File) databases include the transient database, the permanent database, the NCH database, and the security database.

MSAR

Magnetic Storage and Retrieval (MSAR) media provides high speed and high capacity storage libraries on magnetic disk media.

multiple-system committal

Multiple-system committal provides a way to commit copies of documents to more than one system. This may serve as an alternative backup to be used for disaster recovery or as a central repository.

NCH

The NCH (network clearinghouse) is an *MKF* database that keeps track of resources and their addresses. A resource (such as a user ID or a printer) is identified by a three-part name stored in the NCH database in the format object:domain:organization. See object name, domain name, organization name.

network clearinghouse

See NCH.

object

An object is a resource like a tape, printer, database, software service, logon name, etc. See also data object, security object.

object name

The object name is the first part of the *NCH* resource name. Some objects have names predefined by the system. For example, DefaultIMS is the name used to access the index database and you do not change this name. Your service representative configures names for your printers and tape drives.

ODU

An ODU (optical disk unit) is a single optical disk *drive* with no robotic arm. An operator manually inserts disks (media) into the drive and flips the disks over as required.

optical disk

An optical disk is a high-capacity, removable storage device that uses laser technology to store data. Optical disks accept and retain information in the form of marks in a recording layer that an optical beam can read.

optical disk unit

See ODU.

Oracle

Oracle is a relational database created by Oracle Corporation and used for the FileNet Index and WorkFlo Queue databases.

organization name

The organization name is the third part of the NCH resource name (the default name is FileNet). Xerox Corporation registers and distributes organization names.

OSAR

OSAR is an acronym for the Optical Storage and Retrieval unit. See storage library server for a definition of OSAR server functions.

overrides

Overrides cause Security Services to use a security object's attribute instead of the system attributes in certain instances. Overrides apply for these attributes: device security, time use restrictions, security logging, and maximum concurrent sessions.

page cache

Page cache, also known as retrieval cache, is a cache containing all documents being committed to or retrieved from *storage media*. Page cache also stores documents retrieved from media for printing before moving them to print cache. Page cache is an ageable cache.

permanent database

The permanent database stores the media location of each document entered into the system and contains tables for media surfaces, media families, and notes. See database.

permission

Permission is the privilege granted to a *security object* by Security Services to perform certain tasks. Each security object has permissions based on its own definition, as well as the definition of the system and the groups of which the object is a member. Permission is controlled by the extended memberships and override capabilities of the requesting user, groups, and devices.

prefetching

Prefetching is the process of moving a copy of an image from storage media to page cache before a user makes a request for that image. For information about using prefetching, see your client software documentation.

primary family

A collection of media that stores documents from one or more document classes. Before you can complete the definition of a document class, you must assign it to a media family so the system knows where to store documents in that class. You can also create a primary family to be used exclusively for archiving the index and permanent databases.

print cache

System print cache stores image documents waiting to be printed. Application print cache holds all other types of print job data (for example, text reports and files) waiting to be printed. See *cache*.

query

A query is a request for information or the act of requesting information from a database.

radio button

A radio button is a form of input field that is either selected or deselected. Selecting a radio button in a group deselects all the other radio buttons in that group.

RDBMS

RDBMS is an acronym for Relational DataBase Management System. Oracle, DB2 and MS SQL are examples of RDBMSs used in IS. The RDBMS manages the index database and WorkFlo queue database.

remote procedure call

See RPC.

RFS

A RES (Remote Entry Server) is a source Image Services system used for scanning and indexing documents to be committed as a batch on the target Image Services server's storage library.

retention parameters

Retention parameters specify a starting event and a number of months after that event when a document is eligible for deletion. You set up retention parameters when you create a document class.

retrieval

Retrieval is the act of entering a query that results in a list of documents in a query match report. Often, the process includes getting document images from the storage library or document indexing information from the database on the index server.

retrieval cache

See page cache.

retrieval kev

A retrieval key is an index pertaining to certain documents to enable quick document retrieval.

Revise cache

Revise, an optional FileNet product that runs on a PC workstation, is used to manage revisions or change control of technical documents (for example, engineering drawings, blueprints). In Revise, overlays are placed on original documents, marked up, then stored in Revise cache for later viewing or editing where they remain until committed to storage media. See cache.

root/index server

On large systems, you can separate the root and index functions from the storage library functions. The server that checks security, locates devices, and manages the index database is called the root/ index server. See also storage library server.

RPC

RPC (remote procedure call) is a method for a process on one computer to make a request for service on another computer. RPC is the standard throughout the FileNet software for interstation communication.

RSVP message

An RSVP message requires the operator to respond and perform some action, such as loading storage media. An RSVP message can come from any device on the system through the Storage Library Control program.

security database

The security database (an MKF database) contains security information for each object (user, group, device), for the security service, for each direct membership occurrence, and for each function name and class.

security device

The major classes of security devices are terminals and printers and fax servers. Terminal security controls logons and data access from the terminal. Printer and fax server security controls access to print and fax devices.

security logs

Security logs are files into which security event entries are written. Security logs can be viewed through the Events Log option in the Security Administration program.

security object

Each user, group, and device is a security object. Most security characteristics can be applied to users, groups, and devices.

security session

A security session is a single logon occurrence. You can assign users to a session group, which allows you to add or modify logon privileges for an entire group of users at once.

server

A server provides a service in the FileNet system. Types of servers include root and index server, storage library server, and application server.

session

A session is a logical connection allowing two application programs to communicate.

SLAC key

Software License Access Control key. The SLAC key was formerly the software mechanism through which IS controlled access to software services.

SLC

SLC (Storage Library Control) is a FileNet application that monitors the status of storage libraries on the server where SLC is running and provides messages about the media. SLC also provides status information, reports, and controls you use to initiate actions such as enabling or disabling drives, preformatting media, etc. See storage library, storage library server.

SQL

SQL (Structured Query Language) is a standard database query language, pronounced "sequel."

slot

A slot is a storage area for media in a storage library.

station

A station is a hardware component such as a workstation or entry station where a user initiates a request to a server.

storage library

A storage library is a storage media jukebox, a unit that has a number of *slots* for containing storage media and a robotic arm that moves the media between slots, drives, and the input/output slot.

Storage Library Control

See SLC.

storage library server

In a multi-system, the *storage library* server manages the storage libraries and includes cache storage as well as the related databases. A storage library server is sometimes referred to as an OSAR server. A system can have multiple storage library servers. each of which can manage up to eight libraries. In a system with multiple storage library servers, one serves as the document locator server that keeps track of the contents of all storage libraries.

storage media

See media.

System Monitor

An application that displays read-only reports about the state of the FileNet system. The reports are generated from data in the FileNet Management Information Base (MIB), the central database containing Image Services system information.

Task Manager

An application that provides a graphical user interface for controlling and monitoring the FileNet software.

tranlog

See transaction log.

transaction log

The FileNet system is configured to write one or more extra copies of committed documents to transaction log *media*. You can use the predefined media family called tranlog for this purpose, or you can define one or more other media families for transaction logging. The order of images on transaction log media is the order in which images are committed, regardless of the primary media family name and media type

transient database

A directory of documents, images, and available *cache* space. The transient database tracks work in progress, including the batch status, requests to read and write *media*, and print request queues.

user

A security object that can log on to the system and perform tasks.

WAL

The WAL (WorkFlo Application Libraries) is an open, standard interface which allows IBM, Windows, Sun, DEC, and HP workstations to access Image Services and optical disk storage facilities. The interface consists of APIs, accessible through C, C++, Visual Basic, Easel, and PowerBuilder. Referred to more commonly as the Image Services Toolkit (ISTK).

WorkFlo

WorkFlo is the FileNet high-level programming language for imaging processing. WorkFlo is sometimes abbreviated as WFL. Compare with workflow.

workflow

Workflow is the movement of tasks, functions, or activities through a business environment. Compare with WorkFlo.

workstation

A workstation is a client PC that sends requests for services to the server.

WQS

WQS (WorkFlo Queue Services) is the service controlling all WFL queue transactions.

Xapex

Xapex is the Application Executive running in the X-Windows environment.

X-windows

A windowing graphical user interface that runs on the UNIX operating system.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing **IBM** Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing

2-31 Roppongi 3-chome, Minato-ku Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT. MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this

one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation J46A/G4 555 Bailey Avenue San Jose, CA 95141-1003 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be

registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java is a trademark of Sun Microsystems, Inc. in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

U.S. Patents Disclosure

This product incorporates technology covered by one or more of the following patents: U.S. Patent Numbers: 6,094,505; 5,768,416; 5,625,465; 5,369,508; 5,258,855.

Index

| A access rights, document security 102 active documents copying to consolidate media 402 space use report 466 active media | backup logs 50 batch entry services 45 names, system generated 92, 99 security 212 batch cache 53 |
|--|---|
| re-inserting in storage library 441 SLC 449 surface information 459 add (or update) group options 264 administrative group 203 administrator, security 196 ageable cache 51, 53 annotations document copy 397 exporting, DDTS 571 media copy 392 security 212, 546 application print cache 540 server 39 Application Executive 80 application servers 39 | C cache backing up 54 configuration, printing 550 full, printing error 542 location 39 logical 51 objects exported, CSM_exim 475 organization 50 services 45 status, PRI_tool display 511 types 53 cache ageable 51 cache-only systems 40 caching strategy, printing 541 CALS |
| Background Job Control program 365 background jobs controlling 365 deleting event log files 417 monitoring 365 | format 552, 557 product data document 558 cancel print jobs 545 message request 550 requests with PRI_tool 513 |
| tape requests 475 | cataloging 66, 100 |

| CCITT group 4 code, CALS 563 | 2- to 4-digit date 607 |
|---------------------------------|--|
| change | date and time ranges 608 |
| family name, storage media 451 | Windows Server date and time |
| indexes 145 | formats 616 |
| media type, SLC 421 | copying |
| ODU 443 | documents from both sides of media 391 |
| printers, print order 538 | files between servers 501 |
| changing passwords | log of copied documents 392, 397 |
| DB2 355 | requests, disabling storage media 449 |
| character set, support 647 | core files, monitoring 348 |
| character sets, single-byte 662 | cross-system committal 69, 668 |
| choosing a printer or fax 536 | CSM exim |
| clustering 97, 157, 449 | objects included in export 475 |
| COLD | reports 479 |
| documents 71 | running 475 |
| fast batch committal 67 | CSM_tool 50 |
| combined server 38 | cti 487 |
| comments | current |
| DDTS command files 553 | surface, slot and drive map 457 |
| print job header 544 | custom library validation enable/ |
| committal | disable 275, 277 |
| cross-system 69, 668 | custom password validation 273, 274 |
| modes 65 | |
| multiple system 69 | D |
| multiple-system 678 | data |
| to a different system 61 | dictionary, importing/exporting 488 |
| configuration | field, XPR_print 544 |
| printer information 539 | database |
| consolidate media 402 | MKF 48 |
| space use report 466 | monitoring space 525 |
| without deleting documents 402 | overview 47
RDBMS 48–50 |
| Content Federation Services 58 | |
| conventions 31 | database connect administration 353 |
| conversion | database reconnect 362 |
| | Database Server Connect 350 |

| date | timing estimates 556 |
|--------------------------------------|---|
| 2- to 4-digit conversion 607 | transferring multiple documents 558 |
| default formats 612 | ddtsexp command 565 |
| Image Services supported formats 614 | ddtsimp command 575 |
| log file format 602 | declaration file, CALS 558, 559 |
| OS settings 597 | default |
| supported formats 609 | group template 235 |
| valid separators 610 | print cache 541 |
| Windows Server format conversion 616 | user template 236 |
| date mask | delete |
| defining index 136 | group 265 |
| example 607 | index 173 |
| functionality 596 | outdated documents and folders 107 |
| local support 597 | deleted documents, space use report 466 |
| new default 603 | device |
| old default 604 | security 196 |
| DB2 password expiration 356 | type field, XPR_print 544 |
| DDTS 552 | disable |
| abort import or export 554 | cataloging 66 |
| command file 553 | gripper 439 |
| comments 553 | media, storage 420 |
| creating new annotations 581 | optical drive 436 |
| document limit 556 | storage library 433 |
| export 565 | writes only, SLC 449 |
| FileNet options 568 | disk |
| global import options 578 | adding size 116 |
| IDs, document 574 | changing size 117 |
| import 574, 577 | enable/disable drive 436 |
| import document options 580 | display modes 329 |
| options, global 569 | disposition, document class |
| page options 572 | parameters 163 |
| script 583 | distributor queue 68 |
| strings 553 | DMA properties |
| tape support 555 | document class 158 |

| indexes 122 | ODU 443 |
|-----------------------------------|--|
| document | storage library 445, 447 |
| header files 422 | enable |
| ID entry systems compatibility 57 | cataloging, document class |
| location 48 | parameters 159 |
| options, DDTS import 580 | gripper 439 |
| services 46 | media, storage 420 |
| document assembly 63 | optical drive 436 |
| document class 98 | reads only, SLC 449 |
| add new GUID 169 | storage library 433 |
| define DMA properties 158 | entry |
| delete GUID 168 | date, correcting 176 |
| GUIDs 167-171 | systems 57–58 |
| modifying 171 | erasing storage media 399, 402 |
| parameters, description 157 | errors, translating with fn_msg 500 |
| rename GUID 170 | event log |
| report 192 | display modes 329 |
| retention parameters 177 | monitoring 348 |
| security 103, 211 | expiration dates, security 199, 246 |
| document entry 59 | exporting |
| indexing 63 | data dictionary 488 |
| documentation, other 31 | documents, document header file 422 |
| documents | extended membership, security 197, 228 |
| deleting expired 188 | extensible password authentication 272 |
| print cache 540 | extensible user authentication 277 |
| security 211 | F |
| security, printing 546 | f_maint user 350 |
| types 70 | f_open user 350 |
| updating security 185 | f_sqi user 350 |
| domain name 44 | f_sw user 350 |
| drive, add 116 | failed print jobs, priority 545 |
| drop retrieval key 142 | fax server 42 |
| E | fetching to print cache 513 |
| eject media from | file |
| ojoot modia nom | 1110 |

| copies between servers 501 | administrative 203 |
|-------------------------------------|------------------------------------|
| selection 33 | delete 265 |
| system 77 | membership 228 |
| FileNet | primary 204 |
| local directory 33 | rename 267 |
| software 77 | security 195, 201 |
| system definition 37 | update membership 266 |
| filter, query 96 | GUIDs |
| find open documents, BJC report 418 | document class 167-171 |
| fn_msg 500 | user index 125–129 |
| fnlogon 469 | |
| program prompt 471 | H |
| security service 470 | header file, storage media 422 |
| WAL programs, running with 469 | 1 |
| fnsw directory structure 77 | I/O station, storage library 428 |
| folders | identify media 455 |
| security 212 | idle time, printer 547 |
| FolderView tables, spacerpt 525 | Image Services server 38 |
| foreground jobs 364 | importing |
| foreign surface ID 462, 463 | data dictionary 488 |
| function codes | disable storage media 449 |
| Application Executive 586, 591, 592 | documents fail to commit 388 |
| Background Job Control 589 | priority 385 |
| Cache Export/Import Program 590 | procedure 384 |
| COLD Application 591, 592 | reimport selected documents 530 |
| Database Maintenance 587 | selected documents 422 |
| Overriding a Busy Batch 591 | index |
| Server Print Program 586 | counting values 487 |
| Storage Library Control 588 | services 46 |
| future surfaces 467 | storing values on other systems 66 |
| G | index database 64 |
| global options, DDTS import 578 | columns, enlarging 496 |
| grippers, enable/disable 439 | monitoring space 525 |
| group | statistics, ixdb_stat 503 |

| indexes 63 | local surface ID 463 |
|---------------------------------------|--|
| add new GUID 127 | locked cache objects, exporting 477 |
| change retrieval to informational 142 | logical cache 52 |
| date 93 | logical caches 51 |
| define maximum 119 | logon |
| delete GUID 126 | terminate 310 |
| description 122 | times 244 |
| DMA properties 122 | logon database |
| document class parameters 164 | map to Image Services users 287 |
| GUIDs 125–129 | set user's 285 |
| menu 93, 137 | |
| modifying 145 | M |
| numeric 92 | mandatory password change 272, 273 |
| rename GUID 128 | manual mode 429 |
| renaming 147 | manual, conventions 31 |
| string 92, 130 | mask |
| type 124 | default date and time formats 612 |
| verification 166 | time, Image Services supported 615 |
| informational messages, SLC 419, 427 | mask, date 136 |
| insert media in | current functionality 596 |
| ODU 429, 443 | Image Services supported 614 |
| storage library 441 | local support 597 |
| interleaving 88 | new default 603 |
| interrupting printing 545 | old default 604 |
| invalid dates | MASK_OLD 604 |
| correcting 176 | maximum pages per batch, document class |
| J | parameters 161
media 42, 64 |
| - | |
| job header, XPR_print 544 | defining family 109, 115
enable/disable media 420 |
| L | optical, loading 420 |
| less command 509 | size 110 |
| library | space use 466 |
| configuration window 436 | surface summary by library 464 |
| servers 90 | media family |
| | Thedia failily |

| changing during document copy 391, 396 changing name 451 description, SLC report 467 media drives 86 report 193, 467 membership | remote committal 67 global file copy 501 networking 43 remote committal 68, 89, 117 remote committal report 459 retrieval from another system 72 root/index servers 38 |
|--|--|
| adding 263 extended 228 update group 266 | storage library servers 39, 90
multiple-system committal 69, 678 |
| menu index 137 messages, SLC 419, 427 Microsoft SQL Server RDBMS 49 migration 66, 100 delay, document class parameters 160 option, document class parameters 160 to print cache 513 MKF databases 48 modify print request, PRI_tool 516, 551 modifying indexes 145 monitoring batches 335 network activity 332 print jobs 336 multicultural support 647 multiple | N NCH 43 NCH three-part name 43–44 network clearinghouse 43 database 48 network options 42 networking 43 next migrate field, detailed printer status 549 nonactive media, SLC 449 note field, XPR_print 544 notes, security 212 number of copies, XPR_print 544 numeric index increasing values 496 mask 133 |
| copy printing, cache full error 542 storage library servers, print services 541 multiple servers application server 39 changing default domain name 44 cross-system committal 69 dedicated entry systems 57 | O object names 43 security 195 object name 44 obsolete documents and folders 188 ODU 42 eject media 443 |

| insert media 443 | password, custom library validation 275, |
|---|--|
| switch media 443 | 277 |
| open documents, BJC report 418 | password, custom validation 273, 274 |
| operating system security 194 | password, extensible authentication 272 |
| operators 37 | password, mandatory change 272, 273 |
| optical disk drive, enable/disable 436 | password, user expiration exclusion 273 |
| Oracle RAC (Real Application Cluster) 362 | 274 |
| Oracle RDBMS 49 | passwords, changing DB2 355 |
| organization name 44 | pending writes |
| OS settings, date and time 597 | consolidating media 398 |
| outdated media, consolidating 402 | remote committals 459 |
| output, redirecting 473 | performing system management tasks, |
| outstanding print requests 539 | frequency 317 |
| override, security 198 | peripheral devices 41 |
| | perm_DB 39 |
| P8 Content Federation Services 58 | permanent database 39, 48, 61 |
| | permissions, security 196 |
| page cache 53 | prefetches, read requests report 464 |
| pages
batch 161 | preformat media 420, 440 |
| | PRI_tool |
| checking printer status 549
document 161 | hardcopy output 514 |
| printing variable-size 538 | help 515 |
| pager size | idle time 547 |
| modify print request 551 | modify print request 551 |
| XPR_print 544 | next migrate field 549 |
| password change SQL Server 357 | pages printing field 549 |
| password change, DB2 358 | print error field 549 |
| password change, Oracle 357 | redirect print requests 551 |
| password expiration, DB2 356 | primary |
| password failure procedures 357, 358, 359, | group 204 |
| 360, 361 | key 124 |
| password failure procedures, DB2 359, 360, | media, rebuilding 404
primary family 87 |
| 361 | • |
| | assigning 90 |

| print choosing at print time 536 configuration information 539 interrupting job 545 down, troubleshooting 551 job header page note 544 imaging 535 | |
|--|---|
| interrupting job 545 down, troubleshooting 551 | |
| | |
| iob header page note 544 imaging 535 | |
| Jan 1115.515 Paris 2 1 1 | |
| modifying job 544 restarting with PRI_tool 523 | |
| options table 536 selecting with XPR_print 544 | |
| rebooting server 551 status 549 | |
| server for two FileNet systems 540 suspended, restarting 547 | |
| services 535 terminal security 546 | |
| tables in transient database 536 printing | |
| print cache 53 additional reports 544 | |
| configuration options 550 best available 538 | |
| monitoring space 511 multiple storage library servers 54 | 1 |
| strategy 541 network traffic 541 | |
| types 540 troubleshooting 547 | |
| print request UNIX files 542 | |
| canceling with PRI_tool 513 variable-size pages 538 | |
| ID 541 priority | |
| modifying with PRI_tool 516 failed print jobs 545 | |
| outstanding after reboot 539 print requests 538 | |
| priority 538 XPR_print 544 | |
| security 546 | |
| status 521, 536 | |
| print servers 41 | |
| print services 40, 555 | |
| conting with many index 151 | |
| query to coloct a device name 202 | |
| multiple copies 340 | |
| Selecting With AFR_print 544 | |
| print_docs table 556 | |
| print_requests table 536 | |
| print_svcs table 536 raster data file, CALS format 563 printer RDBMS 29, 48 | |

| Microsoft SQL Server 49 | prefetches 464 |
|---|--------------------------------------|
| Oracle 49 | printing 542 |
| read requests | security 309 |
| active media report 459 | surface summary, media 464 |
| by surface, SLC report 464 | reserved status, storage library 436 |
| disabling storage media 449 | resources, changing 45 |
| Real Application Cluster (RAC), Oracle 362 | restart printer 523, 547 |
| reboot print server 551 | retention parameters, updating 176 |
| rebuild storage media 404 | retrieval |
| reconnect, remote database 362 | cache 53 |
| recovery logs 50 | key, defined 95 |
| redirect output, commands 473 | strategies, cti 487 |
| redirected printer, restarting 547 | return code |
| redo logs 50 | menu index 151 |
| reference counts, print cache 541 | root/index server 38 |
| refresh rate, setting for current processes | RSVP messages, SLC 419, 428 |
| table 332 | S |
| reimporting documents that did not | scalar number table 531 |
| commit 530 | scan server 41 |
| remote | scanning 61 |
| committal 68 | security |
| committals, pending 459 | administrative group 203 |
| documents, printing 541 | administrator 196 |
| security service, fnlogon 470 | annotations 212 |
| server, copying files 501 | batch 61, 212 |
| remove media from storage library 447 | changing for committed documents 185 |
| rename | class, document 211 |
| devices 302 | cross-system committal 58 |
| groups 267
users 283 | database 48, 200 |
| reports | device 196 |
| current write surfaces 467 | document 211 |
| find open documents 418 | expiration dates 199, 246 |
| media space use 466 | extended membership 197, 228 |
| modia space use 400 | fnlogon 469, 470 |

| folders 212 | service field, XPR_print 544 |
|---|---|
| group 195, 201 | services 45 |
| logon times 244 | session, security 198 |
| membership group 228 | shared library entry point 275, 278 |
| notes 212 | size of paper, selecting in XPR_print 544 |
| object 195 | SLC |
| options, document class parameters 160 | main window 425 |
| override 198 | on multiple storage library servers 419 |
| override object defaults 245 | tape request messages 475 |
| permissions 196, 254 | space available on storage media 466 |
| planning 252 | ssn 74 |
| primary group 204 | FileNet command 529 |
| reports 309 | foreign media 462 |
| sample setup 256 | standard out 473 |
| server process name and password 305 | statistics, index database 487, 503 |
| session 198 | status |
| system 200 | print requests 535 |
| tabs 212 | printer, PRI_tool 518 |
| template 199 | stdocimp 530–533 |
| terminal and device 295 | stdout 473 |
| user 195, 207 | stopping printer with PRI_tool 513 |
| Security Administration, overview 194 | storage library 39, 42, 426, 428 |
| selecting printer or fax, XPR_print 544 | calibrate 437 |
| serial number, system 529 | enable/disable 433 |
| server | multiple servers running SLC 419 |
| application 39 | removing disabled media 446 |
| combined 38 | storage library servers 39 |
| fax 42 | storage media 42, 64 |
| Image Services 38 | adding to library 441 |
| print 41 | enable/disable 420 |
| process 305 | erasing 399, 402 |
| root/index 38 | insertion guide 423 |
| scan 41 | loading 420 |
| storage library 39 | mapping 455 |

| space available 466 | temporary files, monitoring 348 |
|-------------------------------------|---|
| string index 130 | terminal (printer) security 546 |
| supporting documentation 31 | time |
| surface ID | default formats 612 |
| change to enable/disable media 449 | Image Services supported formats 615 |
| compatibility with entry systems 57 | OS settings 597 |
| ejecting storage media 447 | supported formats 609 |
| local and foreign 462 | valid separators 610 |
| surface information | Windows Server format conversion 616 |
| active media 459 | transaction log 50 |
| summary report 464 | changing type after import 421 |
| switch media, ODU 443 | media 87 |
| system | rebuilding media 404 |
| initialization, print services 539 | rotating media 88 |
| management functions 316 | with interleaving 89 |
| name 43 | transferring files 501 |
| print cache 540, 541 | transient database 48 |
| security 200 | location 39 |
| serial number 529 | print tables 536 |
| setup 73 | |
| System Monitor 315, 322 | U |
| system serial number, ssn 74 | uncommitted images, backing up |
| - | cache 336 |
| T | unified logon |
| tabs, security 212 | delete 83 |
| tape | setup 82 |
| capacity, DDTS 555 | UNIX files, printing 542 |
| request messages, CSM_exim 475 | unknown print server 539 |
| Task Manager 315 | update default |
| access 322
monitor functions 320 | group template options 235 |
| overview 318 | user template options 237 |
| software controls 321 | user security 195, 207
user, expiration exclusion 273, 274 |
| status 320 | user, extensible authentication 277 |
| template security 100 | user, exterisible authentication 277 |

valid separators, date and time 610 verification errors, CSM_exim 475

W

WAL, running with fnlogon 469 WorkFlo queues 68 WorkGroup systems document entry 61 workstations 30 write multiple surfaces 90 number of surfaces 90 write requests active media report 459 by surface, SLC report 464 disabling storage media 449 write surfaces number of 90 report 467

X

XPR_print options 544 using 542-544

IBM.®

Program Number: 5724-R95

Printed in USA

SC19-2689-00

