



Content Services

Windows Oracle Edition

Release 5.5.0

February 2007

| | |
|--|----|
| Getting Started | 10 |
| FileNet Notice to End User - A Software License is Required Prior to Use. | 11 |
| About the Documentation | 12 |
| Documentation Updates | 12 |
| Software Patches | 12 |
| Documentation Formats and Editions | 13 |
| Typographical Conventions | 13 |
| Reader Comments | 14 |
| Trademarks and Copyrights | 15 |
| Software Numbers | 15 |
| Trademarks | 15 |
| Copyright | 15 |
| Addendum | 16 |
| Release Notes | 17 |
| Overview | 18 |
| What is Enterprise Content Management? | 19 |
| Network-wide Search and Retrieval | 19 |
| Logical, Not Physical, Storage Pointers | 19 |
| Object-Oriented Approach | 19 |
| Version Control | 20 |
| Multilevel Security | 20 |
| ECM Administration Tools | 20 |
| Flexible Open Technology | 20 |
| Load Balancing and Scalability of Distributed Services | 20 |
| Data Caching | 21 |
| The Library System | 22 |
| Objects and Properties | 24 |
| Protected Items | 24 |
| Unprotected Items | 24 |
| Document Versions | 25 |
| Users and Groups | 26 |
| Document Storage and Archiving | 28 |
| Storage Repositories and Storage Categories | 28 |
| Archive Repositories and Archive Categories | 28 |
| Searching for Information | 29 |
| Content Searches | 29 |
| Version Tracking | 31 |
| Version Lines | 32 |
| Library System Security | 33 |
| Controlling Access through Groups | 33 |
| How Access Lists Work | 34 |
| Special Privileges for Custom Properties | 36 |
| Secure Document Delete | 37 |
| Clustering | 38 |
| Planning Guide | 39 |
| About the Planning Guide | 40 |
| Who Should Use this Guide | 40 |
| Purpose of this Guide | 40 |

| | |
|---|-----|
| How to Use this Guide | 40 |
| Gathering Information | 41 |
| Whom to Interview | 41 |
| What to Ask | 41 |
| Designing Your Library System | 42 |
| Describe and Diagram Your LAN | 42 |
| Estimate Hardware and Software Requirements | 42 |
| Configuration Worksheets | 43 |
| Users and Groups Worksheets | 45 |
| Administrators Worksheet | 47 |
| Document Description Worksheet | 48 |
| Document Access Rights Worksheet | 49 |
| Network Worksheet | 51 |
| Budgetary Constraints Worksheet | 52 |
| Document Storage & Access Worksheet | 53 |
| Site Hours Worksheet | 54 |
| Network Usage Worksheet | 55 |
| Document Timeframes Worksheet | 56 |
| LAN Description Worksheet | 57 |
| Machine Requirements Worksheet | 58 |
| Library System Components and Sizing Worksheet | 59 |
| Client Application Requirements Worksheet | 61 |
| Library System Requirements Worksheet | 62 |
| Database Server Hardware Worksheet | 64 |
| Initial Storage Server Hardware Worksheet | 65 |
| Additional Storage Server Hardware Worksheet | 66 |
| Disaster Recovery Worksheet | 67 |
| Roadmap to Setting Up | 68 |
| Installation Roadmap | 69 |
| Administrative Client Setup Roadmap | 71 |
| Configuration Roadmap | 72 |
| Client Interface Setup Roadmap | 74 |
| Upgrade Roadmap | 75 |
| Installation Guide | 76 |
| Installation Overview | 77 |
| Installation Constraints | 77 |
| Installation Scenario | 78 |
| Installation Platforms | 79 |
| Prerequisites to Installing Content Services | 80 |
| Installing Content Services with Image Services | 80 |
| Installing Content Services into Non-English Environments | 81 |
| Windows 2000 or Windows 2003 Requirements | 81 |
| Server Name Requirements | 81 |
| Database Name Requirements | 82 |
| Library System Name Requirements | 82 |
| Content Services in a Cluster Environment | 83 |
| Setting Up the Oracle Software | 89 |
| Optimizing Oracle | 108 |
| Connecting to the Oracle Database | 109 |
| Creating a Database User | 115 |
| Procedure for Moving Indexes to a Separate Tablespace | 117 |
| Configuring for WANs and Firewalls | 117 |
| Installing Content Services | 119 |

| | |
|---|-----|
| Installation Settings | 119 |
| Running Content Services Setup | 123 |
| Installing Initial Library System Services | 123 |
| Configuring a Library System Within a Cluster | 127 |
| Starting and Monitoring Library Services | 128 |
| What to Do Next | 129 |
| Accessing and Administering a Library System | 130 |
| Installing Additional Library System Services | 130 |
| Adding a Display Language | 134 |
| Installing Content Services Unattended | 136 |
| Upgrading the Oracle Database | 143 |
| Upgrading to Oracle 10g 10.2.0.1.0 on Windows | 144 |
| Upgrading to Oracle 10g 10.2.0.1.0 on UNIX | 146 |
| Exporting/Importing a Windows Oracle Database | 148 |
| Exporting/Importing a UNIX Oracle Database | 150 |
| Upgrading Content Services | 152 |
| Upgrade Overview | 152 |
| Upgrade Constraints | 153 |
| Preparing for the Upgrade | 154 |
| Preparing Verity K2 for the Upgrade | 156 |
| Upgrading Library System Services | 157 |
| Rebuilding Content Search Repositories | 160 |
| Upgrading Content Services Unattended | 160 |
| Separating a Database from Library Services | 163 |
| Configuring Content Services | 164 |
| The CS Configuration Tool | 164 |
| Understanding Default Values and the Inheritance Path | 188 |
| Setting Default Properties | 190 |
| Identifying the Primary Storage Manager | 191 |
| Setting Version Online Limits and Counts | 191 |
| Setting Version Index Limits and Counts | 193 |
| Controlling Search Results | 193 |
| Minimum Item Delete Access Level | 195 |
| Customizing Logon Procedures | 195 |
| Assigning a Display Language | 196 |
| Changing Operating System and Database Display Language | 196 |
| Controlling Library System Services | 198 |
| Content Services Manager | 198 |
| The Control Panel Services Tool | 201 |
| Controlling Library System Services with Command-line Utilities | 202 |
| Controlling Library System Services in a Cluster Environment | 202 |
| Accessing Multiple Library Systems | 204 |
| Multiple Library Systems With One Database Engine | 204 |
| Adding a Library System | 205 |
| Connecting a Client to Different Library System Platforms | 206 |
| Client Library Files | 207 |
| Prerequisites to Installing Client Library Files | 207 |
| Installing Client Library Files | 208 |
| Installing Client Library Files Unattended | 208 |
| Client Library Files Setup Output File | 211 |
| Client Library Files Setup Error Messages | 212 |
| All UNIX Clients | 214 |
| Problem Solving | 215 |
| Uninstalling Content Services | 217 |
| Uninstalling Library Services | 217 |
| Uninstalling Verity Software | 218 |
| Uninstalling Library Services Unattended | 218 |

| | |
|--|------------|
| Uninstalling Content Services in a Cluster Environment | 219 |
| Uninstalling Client Library Files | 219 |
| Uninstalling Client Library Files Unattended | 220 |
| Uninstalling Client Library Files via MS-SMS | 220 |
| Administration Guide | 222 |
| Configuration And Maintenance | 224 |
| Storage Space Overflow | 224 |
| Avoiding Property Repository Overflow | 225 |
| Managing Storage Space | 226 |
| Preparing for a Possible Repository Failure | 227 |
| Adjusting Storage Manager Configurations | 228 |
| Configuring Databases | 229 |
| Setting the Database Connections | 229 |
| Setting Memory Allocation | 230 |
| Setting the Redo Logs Buffer Allocation | 231 |
| Setting Database Buffers Allocation | 232 |
| Storage Repositories | 233 |
| Adding a Storage Repository | 234 |
| Modifying a Storage Repository | 235 |
| Moving a Storage Repository | 235 |
| Moving Versions Between Storage Repositories | 236 |
| Setting Up a Direct-Document-Access Repository | 237 |
| Automatic Management of Storage Repositories | 238 |
| Storage Categories | 240 |
| Adding a Storage Category | 241 |
| Modifying a Storage Category | 241 |
| Archive Repositories | 242 |
| Moving Versions Offline | 243 |
| Reclaiming Offline Versions | 244 |
| Processing a User's Reclaim Request | 245 |
| Adding an Archive Repository | 247 |
| Modifying an Archive Repository | 248 |
| Moving an Archive Repository | 248 |
| Archive Categories | 250 |
| Adding an Archive Category | 251 |
| Modifying Archive Categories | 251 |
| Content Search Repositories | 253 |
| Adding a Content Search Repository | 255 |
| Modifying a Content Search Repository | 256 |
| Moving a Content Search Repository | 256 |
| Custom Properties | 258 |
| Internal Custom Properties | 258 |
| Characteristics of a Custom Property | 258 |
| Adding Custom Properties | 259 |
| Modifying Custom Properties | 261 |
| User Objects | 263 |
| Working with User Objects | 263 |
| Assigning and Modifying Values in a User Object | 264 |
| Group Objects | 267 |
| Working with Group Objects | 268 |
| Assigning and Modifying Values in a Group Object | 269 |
| Special Administrative Groups | 271 |
| Importing Users and Groups | 272 |
| Using Windows User Import | 273 |
| Installing an iPlanet LDAP Server | 276 |

| | |
|---|-----|
| Using LDAP User Import | 277 |
| Using XML User Import | 281 |
| Passwords for Imported Users | 284 |
| User Import Usage Guidelines | 285 |
| User/Group Synchronization Criteria | 292 |
| Windows User Import Usage Example | 294 |
| Windows User Import Log File Example | 296 |
| Sample LDAP Export XML File | 297 |
| LDAP Export XML Schema | 300 |
| User Import Preview Transaction File | 307 |
| Managing Multiple Library Systems | 310 |
| Sharing Servers and Disk Space | 310 |
| Optimizing Search Queries | 312 |
| Using Indexed Properties | 312 |
| The LIKE Operator | 314 |
| Using Secure Searches | 314 |
| Specifying Keywords in Searches | 314 |
| Using Content Search Criteria | 314 |
| Content Search Filters | 315 |
| Configuring Content Search Stopwords | 320 |
| Function-based Indexes | 320 |
| Audit Logging | 323 |
| Audit Logging Features | 323 |
| Setting Up and Managing Auditing | 324 |
| Audit Log File Format | 325 |
| Managing the Audit Log File | 327 |
| Audit Log Reports | 329 |
| Secure Document Delete | 331 |
| Using Secure Document Delete | 331 |
| Secure Document Delete Limitations | 333 |
| Deleting Sessions | 334 |
| Resetting the Session ID | 334 |
| Backing Up the Library System | 336 |
| Planning the Backup | 336 |
| Housekeeping | 337 |
| Backing Up Property Repositories | 337 |
| Backing Up Storage Repositories | 337 |
| Backing Up Archive, Reclaim, and Reclaim Command Repositories | 337 |
| Backing Up Content Search Repositories | 338 |
| Restoring Content Services | 339 |
| Restoring the Database | 339 |
| Restoring Library Services | 340 |
| Restoring Storage Repositories | 343 |
| Restoring Archive Repositories | 345 |
| Restoring Content Search Repositories | 346 |
| Renaming Servers | 348 |
| Moving Content Services | 352 |
| Moving the Database | 352 |
| Moving Library System Services | 355 |
| Moving a Verity K2 Master Admin | 360 |
| Reinstalling CS Executable Files Unattended | 364 |
| Document Classes and CVLs | 366 |
| System Messages | 367 |
| Suspending a Storage Manager | 368 |
| Utilities | 369 |
| ARCHDATE | 371 |

| | |
|---|------------|
| CAT_INDX | 375 |
| CAT_PROC | 377 |
| CAT_STAT | 379 |
| CCAT | 380 |
| COPY_CP | 382 |
| CSADMIN | 384 |
| CSK2CFG | 386 |
| CS_IMPORT | 392 |
| DMU | 394 |
| DSINFO | 398 |
| DSSTOP | 400 |
| DSSTRT | 402 |
| DTCVLCHK | 405 |
| DTCVLMIG | 412 |
| FILE2ID | 419 |
| HMSRVCHK | 420 |
| HMSRVMIG | 424 |
| ID2FILE | 428 |
| INDXDATE | 429 |
| INDXFIX | 432 |
| LDAP_EXP | 435 |
| NLSADMIN | 438 |
| REPCHECK | 439 |
| REPLKICK | 443 |
| SPISTAT | 445 |
| SRDELETE | 446 |
| USGRPCHK | 448 |
| USGRPMIG | 455 |
| VUPDATE | 461 |
| Creating Reports Using Database Schemas | 462 |
| Changing Existing Reports | 463 |
| Database Schema | 464 |
| Sample SQL Queries | 569 |
| Status Codes | 573 |
| Status Codes, Return Values, and Descriptions | 573 |
| Replication Services | 593 |
| Replication Services Overview | 594 |
| Replication Services Terminology | 595 |
| Objects Supported for Replication | 597 |
| Administrative Domains | 598 |
| Outside an Administrative Domain | 599 |
| The Replication Daemon | 599 |
| The Replication Home Server feature | 599 |
| Managing Document Classes and CVLs | 600 |
| Planning for Replication Services | 601 |
| The Scope of Replication Services Needed | 601 |
| Existing Library Systems | 602 |
| The Intended Distribution of Replicated Documents | 602 |
| Impact on Network and Hardware Resources | 603 |
| Backup and Restore Operations | 603 |
| The Original Library System | 604 |
| Administrative Domains | 606 |
| Replicating Objects | 609 |
| Replicating Documents and Folders | 609 |
| Replicating Compound Documents | 612 |

| | |
|--|------------|
| Replicating Published Documents | 612 |
| Replicating Document Annotations | 612 |
| Replicating Document Classes and CVLs | 612 |
| Replicating Custom Property Values | 613 |
| Replicating Users and Groups | 614 |
| Replicating Stored Searches | 614 |
| Replicating Document Item and Version Properties | 615 |
| Replication Home Server | 616 |
| Centralized management of documents and folders | 616 |
| Adding subfolders and documents on an offline Domain Master | 617 |
| Applicable only to member libraries within the administrative domain | 617 |
| Setting Up Replication Services | 618 |
| Allowing for Replication Latency | 619 |
| Configuration Results | 620 |
| Editing the tnsnames.ora File for Replication Services | 621 |
| Establishing a Replication Administrative Domain | 621 |
| Configuring Replication Folders | 623 |
| Replication Services Configuration Parameters | 624 |
| Enabling Replication Home Server | 625 |
| Using the Folder Administration window | 625 |
| Using the HMSRVMIG utility | 626 |
| Exceptions to migration to Home Server | 628 |
| Managing Replication Services | 629 |
| Managing Replication Folders | 629 |
| Modifying Document or Folder Replicas | 632 |
| Managing Users and Groups | 633 |
| Managing Documents | 634 |
| Starting, Stopping, and Disabling Replication Services | 636 |
| Replication Guidelines for End Users | 637 |
| | 641 |
| Managing Replication Home Server | 644 |
| Operations allowed on a Home Server enabled replication folder | 644 |
| Disabling Replication Home Server | 644 |
| Replication Recovery | 645 |
| Working of the Replication Recovery Utility | 646 |
| Preparing for Replication Recovery | 648 |
| Running the Replication Recovery Utility | 650 |
| Replication Recovery Utility Limitations | 651 |
| Replication Planning Worksheets | 652 |
| Replication Users and Groups Worksheets | 653 |
| Replication Document Class and CVL Worksheet | 655 |
| Administrative Domains Worksheet | 657 |
| Rendition Services | 658 |
| Rendition Services Components | 659 |
| Hardware and Software Requirements for Rendition Services | 659 |
| Uninstall Previous Versions of Rendition Services | 660 |
| Installation Overview | 660 |
| Prepare a Server | 662 |
| Install Adobe Acrobat Distiller Server 5.0.5 | 663 |
| Install Rendition Services on Servers | 664 |
| Rendition Services Install | 664 |
| Configure the Rendition Services Server | 665 |
| Complete the Server Setup | 671 |
| Multi-Server Install | 675 |
| Disable Options in Microsoft Word and Corel WordPerfect | 677 |

| | |
|---|-------------------|
| Prepare Microsoft Word Documents | 677 |
| Prepare Corel WordPerfect Documents | 678 |
| Configure the Corel WordPerfect Printer | 678 |
| Optional Steps | 680 |
| Configure Server Content Types | 680 |
| Application | Configuration 680 |
| | 681 |
| Access Processing Preferences for an Application | 681 |
| | 681 |
| Associate Plug-ins with Content Types | 681 |
| Set Plug-in Options | 682 |
| Plug-in Options and Processing Preferences | 683 |
| Default and Available Plug-in Options | 696 |
| Remap Special Characters | 698 |
| Font Character Sets | 700 |
| Check the Status of Rendition Services Jobs | 703 |
| Domain Configuration from Vista Administration | 703 |
| Configure Content Types from Vista Administration | 703 |
| Work with Rendition Services Jobs | 704 |
| Control Rendition Services Jobs | 706 |
| | 706 |
| RenderPerfect | 706 |
| Action Groups | 707 |
| | 710 |
| Folder Groups | 710 |
| | 712 |
| Managing Folder Groups | 712 |
| Watch Items | 712 |
| Specify Archived Jobs | 714 |
| Watched Folders and IDM Desktop 4.0 (or higher) Preferences | 714 |
| Log Files | 715 |
| Uninstall Rendition Services on a Server | 717 |
| | |
| Glossary | 718 |
| General Terms and Concepts | 718 |
| Replication-specific Terms and Concepts | 720 |
| | |
| Index | 722 |

Getting Started

| For information on: | Refer to: |
|---|---|
| Late Breaking News | “Release Notes” on page 17 FileNet Worldwide Customer Support |
| Content Services | “Release Notes” on page 17 “Overview” on page 18 |
| Documentation | “About the Documentation” on page 12 |
| Installing Content Services | “Roadmap to Setting Up” on page 68 “Planning Guide” on page 39 “Prerequisites to Installing Content Services” on page 80 “Installing Content Services” on page 119 |
| Administering Content Services | “Administration Guide” on page 222 |
| Administering Replication Services | “Replication Services” on page 593 |
| Installing, Configuring, and Administering Rendition Services | “Rendition Services” on page 658 |

FileNet Notice to End User - A Software License is Required Prior to Use.

BEFORE COMPLETING INSTALLATION OR USING THIS SOFTWARE, CAREFULLY READ THIS NOTICE. THIS SOFTWARE IS THE PROPRIETARY INTELLECTUAL PROPERTY OF FILENET CORPORATION (OR ITS SOFTWARE SUPPLIERS) AND USE OF ANY PORTION OF THE SOFTWARE IS ONLY PERMITTED IF YOU HAVE A VALID WRITTEN LICENSE AGREEMENT WITH FILENET.

You may have a valid FileNet software license agreement if:

1. Your EMPLOYER and FileNet have entered into a written license agreement; or
2. Your EMPLOYER and an authorized FileNet partner have entered into a written license agreement.

If you do not have a valid license agreement to use the software, then terminate the installation of this software, promptly delete any FileNet software files from your computer, and return the software media and all other related items to: FileNet Corporation, 3565 Harbor Blvd., Costa Mesa, CA 92626-1420, USA. FileNet only licenses use of Software to end user. Nothing in this Notice or any written license agreement constitutes a sale of the software. FileNet reserves all rights.

Notice to U.S. Government of Restricted Rights.

The software and documentation are "Commercial Items," as that term is defined at 48 C.F.R. §2.101 (October 1995), consisting of "Commercial Computer Software" and "Commercial Computer Software Documentation," as such terms are used in 48 C.F.R. §12.212 (September 1995) or 48 C.F.R. §227.7202-1 through 7202-4 (June 1995) as applicable.

Consistent with 48 C.F.R. §12.212 or 48 C.F.R. §§227.7202-1 through 7202-4, as applicable, the Commercial Computer Software and Commercial Computer Software Documentation are being licensed to U.S. Government end users (a) only as Commercial Items, and (b) with only those rights as are granted to all other end users pursuant to the terms and conditions of the applicable software license agreement.

Contractor/manufacturer is: FileNet Corporation, 3565 Harbor Blvd., Costa Mesa, California 92626-1420, USA.

Unpublished-rights reserved under the copyright laws of the United States.

About the Documentation

The following two procedures show how you can update your Content Services documentation and obtain Content Services software patches.

Documentation Updates

To view or download an update of Content Services documentation:

1. Go to the [FileNet Worldwide Customer Support](#) web site.
2. If you have a user name and password, enter them on the Welcome to Worldwide Customer Support page, and click Login. Then go to [step 4](#). Please note that your user name and password are case-sensitive. If you require assistance, click **Contact Us** at the bottom of the web page.
3. If you do not yet have a user name and password, click the New Users **sign up** link. As a new user, you need to ensure that the following conditions have been met before attempting to log on to the Web site:
 - Your company has a support contract with FileNet CSS.
 - You are an authorized customer contact for this account.
 - You have your FileNet Customer ID (CID), which you can find on FileNet correspondence, such as packing lists and invoices.
4. On the Customer Service & Support page, click **Product Tech Info**.
5. In the Products list on the left side of the Products page, click **Image Manager (IM)**.
6. In the Products list on the left side of the Products page, click **Content Services**.
7. In the Products list in the center of the Products page, click **Product Documentation**.
8. Scroll the Products list in the center of the Products page to the product whose updated manual, help file, or readme file you want and click the link.
9. Scroll the expanded Products list for the link to the documentation file you want to access.
10. Click the link to open the document, or right-click and choose Save Target As to download the file.

Software Patches

To download a patch:

1. Go to the [FileNet Worldwide Customer Support](#) web site.
2. If you have a user name and password, enter them on the Welcome to Worldwide Customer Support page, and click Login. Then go to [step 4](#). Please note that your user name and password are case-sensitive. If you require assistance, click **Contact Us** at the bottom of the web page.
3. If you do not yet have a user name and password, click the New Users **sign up** link. As a new user, you need to ensure that the following conditions have been met before attempting to log in to the Web site:

- Your company has a support contract with FileNet CSS.
 - You are an authorized customer contact for this account.
 - You have your FileNet Customer ID (CID), which you can find on FileNet correspondence, such as packing lists and invoices.
4. In the Service Window, under Download Patches, click **FileNet Products**.
 5. On the FTP directory page, click **IDMDocServices.html** to open the Content Services FTP patch download directory.
 6. Click the directory link for the software release number (**5.5x**, for example) to open the list of available hotfix packs and patches.
 7. Click the directory link for the hotfix pack or patch you want to download.
 8. Read the available readme file and download the software ZIP file as instructed.

Documentation Formats and Editions

The Content Services documentation is provided as Portable Document Format (**cs*.pdf**), optimized for viewing with Adobe Acrobat Reader 4.0 or later. These files are provided for users who wish to print the documentation in book format. These files are created with compatibility set for Acrobat 4.x. There are known problems (such as some characters and graphics dropping out) associated with viewing them using earlier versions of the Acrobat Reader.

Note The asterisks (*) above indicate that a platform-specific abbreviation completes the file name for each edition of the Content Services documentation provided:

- csWinSql = Windows SQL Server
- csWinOracle = Windows Oracle
- csHpOracle = HP-UX Oracle
- csSolOracle = Solaris Oracle

Typographical Conventions

The Content Services documentation uses these conventions to distinguish elements of text.

| Convention | Usage |
|---------------|--|
| UPPERCASE | Environment variables, status codes, utility names. |
| Bold | Paths and file names, program names, and selected terms such as command parameters or environment variables that require emphasis. |
| <i>Italic</i> | User-supplied variables and new terms introduced in text. |
| Monospace | Code samples, examples, display text, and error messages. |

Reader Comments

If you have comments, questions, or suggestions regarding this documentation, we in FileNet Documentation will be happy to respond. Please write or call:

FileNet Corporation
720 4th Avenue, Suite 100
Kirkland, Washington 98033
(425) 893-7000

Trademarks and Copyrights

Software Numbers

Software Release: 5.5

Information in this document is subject to change without notice. The software described herein is furnished under license and maintenance agreements, and it may be used or copied only in accordance with the terms of the agreements. No part of this manual may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage or retrieval systems, for any purpose other than the purchaser's personal use without the express written permission of FileNet Corporation.

Trademarks

FileNet is a trademark of FileNet Corporation. Microsoft and Windows are registered trademarks and Windows NT, Windows 95, Windows 98, Windows XP, and Internet Explorer are trademarks of Microsoft Corporation. Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation. UNIX is a registered trademark of X/Open Company Limited. Oracle is a registered trademark of Oracle Corporation. Solaris is a registered trademark of Sun Microsystems, Inc. HP-UX is a trademark of Hewlett-Packard Company. All other products mentioned are trademarks or registered trademarks of their respective manufacturers.

Copyright

© 2001, 2007 FileNet Corporation. All rights reserved.

Portions © 2006 Verity, Inc. All Rights Reserved.

© 2003 Liquent, Inc. All rights reserved. This software is provided under a license agreement with Liquent. Use of this software and all related documentation is subject to and limited by the terms and conditions stated in such software license agreement.

Addendum

A horizontal dotted line in teal extends from the left margin to the right edge of the page. At the right edge, it turns 90 degrees downward, continuing as a vertical dotted line. A solid teal rectangular block is positioned in the top right corner of the page, partially overlapping the horizontal dotted line.

Any notes in this section describe updates made to the documentation since the Content Services 5.5 release.

Note For details on how to download updated or auxiliary documentation and software patches related to Content Services, see [“About the Documentation”](#) on page 12.

Release Notes

A horizontal dotted line in teal color extends from the left margin towards the right. At the point where it meets a solid teal vertical bar on the right side of the page, the line turns 90 degrees downwards, continuing as a vertical dotted line.

CS 5.5.0 Windows Oracle Release Notes, accessible at [FileNet Worldwide Customer Support](#), contains three topics:

- **Release Highlights** describes what is new in CS 5.5.0 or changed since CS 5.4.0.
- **Problems Fixed in This Release** contains a list of the bugs, with Defect numbers included for reference, found in earlier versions of Content Services that have been fixed in CS 5.5.0.
- **Known Issues** contains a list of known problems and critical information for CS 5.5.0, with Defect Numbers included for reference.

Refer to “[About the Documentation](#)” on page 12 for instructions on how to access web-posted Content Services documentation.

Overview

| For information on: | Refer to: |
|--|--|
| The basics of Enterprise Content Management and Content Services | “What is Enterprise Content Management?” on page 19 “The Library System” on page 22 |
| How objects and properties are used in a library system | “Objects and Properties” on page 24 |
| How Content Services stores and archives documents in a library system | “Document Storage and Archiving” on page 28 |
| The search features available for finding information in a library system | “Searching for Information” on page 29 |
| How versions of the items checked into a library system are managed | “Version Tracking” on page 31 |
| How documents stored in a library system are protected against unauthorized access | “Library System Security” on page 33 |
| How Content Services is configured in a cluster environment | “Clustering” on page 38 |

What is Enterprise Content Management?

Corporate information comes in many forms, from print to voice to electronic. An enterprise content management (ECM) system spans an entire enterprise, from site to site, across LANs and WANs, enabling all the people in your organization to access, protect, share, and manage corporate information.

An ECM solution generally includes:

- Network-wide search and retrieval
- Logical, not physical, references to storage locations
- Object-oriented approach
- Version control
- Multilevel security
- Network administration tools
- Flexible and open technology
- Load balancing and scalability of distributed services
- Data caching

Network-wide Search and Retrieval

FileNet Content Services (CS) systems allow documents to be identified using plain language titles and property names, not cumbersome file names.

End users can search for documents based on their properties and/or their content. Regardless of the user's location or the file's location, complex path commands with coded document names are never required. Documents are checked out of and checked in to the FileNet CS system for updates and revisions. One simple end user command transparently saves, tracks, and archives any document.

Logical, Not Physical, Storage Pointers

Unlike conventional document management packages that rely on physical file locations, a FileNet CS system does not require users to have any knowledge of where documents are actually stored. In fact, users are shielded from that knowledge, providing an added layer of security for the documents and the network.

Document management systems that use physical pointers require constant supervision and cumbersome revision as networks grow and change. Moving files, adding users, adding or changing directories means updating every single workstation individually. With a FileNet CS system, such system changes can be handled centrally by the network administrator.

Object-Oriented Approach

A FileNet CS system encapsulates all types of files and objects. Each object is described by properties such as description, security, and administration. Actions are performed against objects using these properties as variables. In this way, an ordinary file becomes an “intelligent” document because, regardless of how or where it is used, the

document retains its properties and always knows who is authorized to view it or make edits, where it is to be stored on the network, and how it is to be administered and archived.

Version Control

When a user opens a document for editing, it is designated as checked out. The user is always provided a copy of the particular version of the document, and the original copy of that version remains protected. Other users who try to access a document version that is checked out are notified that it is checked out and by whom. While the document is checked out and being modified, users can view the latest checked-in version of the document.

When the user is finished working with a checked out version, he or she can check it in. The FileNet CS system retains the old version and automatically saves the checked in document as a new version. Once checked in, the old and new versions of the document become available for check out to other users. When users access the document, they are free to check out the most current version or go back to older versions as needed.

Multilevel Security

An access control list is created for each document added to a FileNet CS system. This access control is part of the document's properties and is independent of the document's location on the network. In this way, authorized users have access to the information they need, no matter where they or the documents are physically located. Yet all stored documents are well protected, because only the documents that each user is authorized to see will be displayed as the result of a search.

A FileNet CS system controls who can create new versions of any given document, regardless of data type, application used, or storage location. Security levels determine access rights such as *Owner*, *Author*, *Viewer*, or *None*. The only access to documents is through the FileNet CS system. File names are encrypted and users are prevented from bypassing the FileNet CS system to access files via the network operating system. CS administrators should not share the directories where FileNet file are located.

ECM Administration Tools

FileNet CS systems include a Windows[®] based administrative tool. ECM administration tasks such as adding new servers, designating new storage locations, or adding new users and groups can be performed at any time, no matter how many users are online. In addition, a web browser-based version of the administrative tool that includes a subset of the windows-based administration tool is available for administrators who want mobile access to library system properties and the ability to start/stop CS servers remotely.

Flexible Open Technology

CS is designed to work with virtually every combination of user interface, application, and hardware/software. Creation and maintenance of general office documents (including enterprise, departmental, workgroup, and personal documents), legal document management, email, as well as document imaging may all be handled through a common repository. This means that no matter how many platforms or environments are used across the network, all enterprise applications are interoperable. Additionally, any enterprise can take advantage of FileNet IDM Desktop (and, if necessary, CS API programming functions) to build custom enterprise applications as future needs emerge.

Load Balancing and Scalability of Distributed Services

No matter how an organization's workload grows and changes across departments or divisions, the FileNet CS system's queue management and load balancing features ensure a consistent working environment for end users.

All user requests for documents are processed through efficient multithreaded or multiprocess services, allowing simultaneous requests to be processed in parallel. If a server is unavailable, the request is automatically handled by another server. FileNet CS systems can be configured to evenly allocate document storage as well. If the traffic load becomes too great, additional servers can be added to improve system performance for end users.

Data Caching

The data caching features of CS speed up access to information that is continually needed internally by application programs and stored procedures. For example, when a user opens a CS session (thereby creating a Session object), CS caches the session information. Thereafter, the application programs can use the information without having to access the database after logon has occurred.

Caching also speeds up the assignment of initial property values for many newly created objects, such as Item objects. For example, when a user adds a new Item, the initial default Item property values are based on information stored in the session context. This information is accessed (and cached) when the user session begins, not when the Item object is created.

To reduce the overhead required to process searches on the CS server, search results are cached on the client machine, and rows are retrieved locally.

Objects that are retrieved from the database are cached in a local object allocation buffer (OAB) for use by application programs for the lifetime of the object, thereby caching intermediate SQL results from the database in addition to the search and session context data.


The Library System

Content Services (also called the *library system*) provides server-based enterprise content management (ECM) that can be accessed via a client interface using applications such as FileNet IDM Desktop or FileNet Web Services. An enterprise, especially if it's spread across various sites, may have multiple library systems, depending upon sizing and load balancing considerations.

A typical Content Services installation comprises the following components:

- One or more library systems.
- One or more administrative clients.
- One or more ECM clients (such as IDM Desktop).
- One or more application solutions (such as FileNet Web Services for intranet/internet access).

As you can see in the following illustration, each library system has a set of sub-components that you install and configure to provide the services for your various client-based users.



A Basic Library System

Each library system acts as an information storehouse. The Property Manager stores and manages object properties (users, groups, system, and custom properties) and metadata describing your cataloged documents. The Storage Manager stores the actual documents and process requests. The Content Search Manager provides document

content-based indexing and searching capabilities. In addition, you have the option to enable CS Replication Services, which allows you to replicate documents across multiple library systems.

Note You have the option to locate your database on a standalone property server. In this configuration, the initial Storage Manager, Content Search Manager, and Replication Services would be located not on the property server, but on the initial storage server.

Objects and Properties

A library system is based on an object-oriented architecture: users and administrators add instances of objects to the library system which then manages these objects. Objects represent real things, such as electronic files, hard-copy documents, drawings, users, groups, and library system components. Although a library system manages many different types of objects, those that users will work with most frequently are:

- Items
- Versions
- Users
- Groups

Of these objects, users will add only items and versions to the library system. Although as the administrator, you may give users the access privileges to add and modify User and Group objects, you usually take care of adding users, groups, and many other objects that will help you manage the library system and help your users manage their documents.

When something such as a document or user is added to a library system, the system creates an instance of the related object for it (in this case, the Item or User object) and catalogs the properties describing this object in an electronic property repository.

The library system automatically assigns values for some properties. For example, when you add a user to the library system, a value is automatically assigned to the Date Added property. There are other property values of the User object that you can assign when you add the user to the library system, such as the values you want to assign to the User Name and Comment properties.

Some properties require values when you add an instance of an object. For example, when you add a user to a library system, you must assign a value for the User Name property in the User object.

For Item, Version, and Checkout objects only, there are properties called custom properties, which you can define yourself. For example, if a particular library system has been installed in a word-processing environment, you might create custom properties such as Subject, Author, Work Order Number, and Supervisor for your Item or Version objects.

Protected Items

When you add a document to a library system, you usually add it as a protected item. It is called “protected” because the library system physically stores the document, thereby protecting it from unauthorized access, inadvertent removal, and more.

Protected items can be documents that were created by any application (for example, word-processing documents, financial spreadsheets, images, audio, tables, and so on).

Unprotected Items

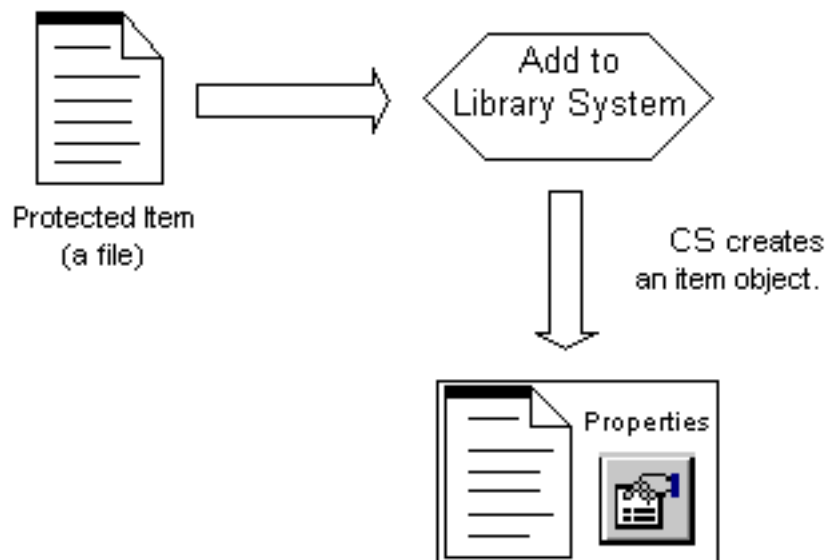
A library system also lets you control and manage unprotected items (external documents), which are items that are not physically stored in the library system but are tracked for easy management. This capability is especially useful

for items that are stored in special locations, yet still require version tracking or controlled access. Some examples are: magnetic tapes, printed technical manuals, printed engineering drawings, and software source code listings.

Although the unprotected items added at most sites represent these kinds of physical objects, unprotected items can also be electronic files that you do not want to store in the library system.

Document Versions

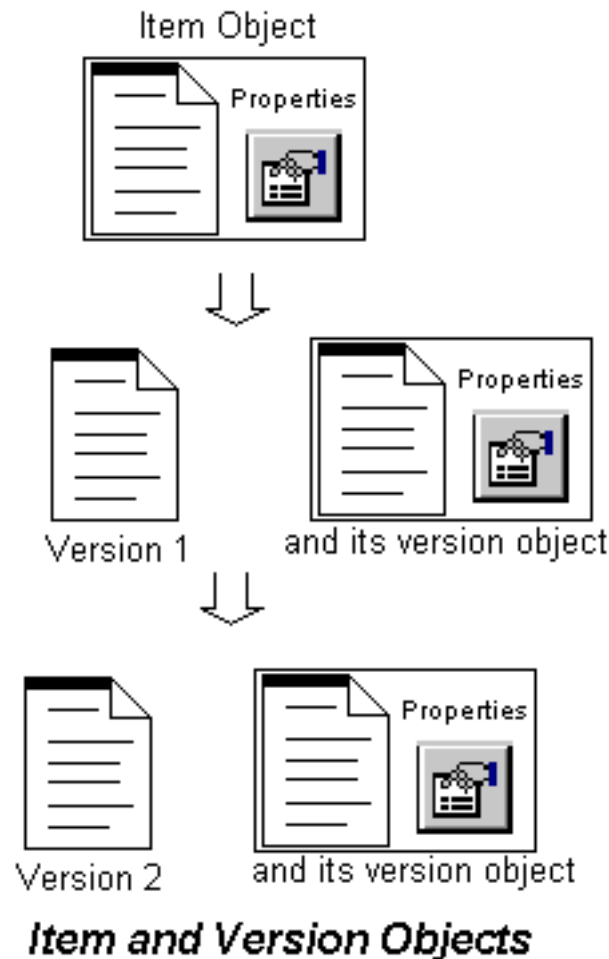
When a document is added to a library system, an instance of the Item object is created to describe its general features or properties: the author, the date it was added, the original file name, and so on.



Adding an Item to a Library System

At the same time that an instance of the Item object is created to generally describe the document, an instance of the Version object is created to uniquely identify the first version of the document. A version is always associated with a document. This Version object includes information such as the user who created the version, the date the version was created, and any comments about what makes this version unique.

A Version object acquires more meaning as revisions are made to the document. If you were to modify version 1, for example, you would be expected to save the new revision as version 2 of the document. As a result, a second instance of the Version object would be created to describe who made the changes, when the changes were made, and any comments you might have.



Each time you create a new version of a document, the library system creates a new instance of the Version object to uniquely identify any changes. Thus, each document has one instance of the Item object to describe the original document and, possibly, several instances of the Version object, each one identifying a revision of the document.

Users and Groups

Before anyone can access the information managed by a library system, you must first add that person to the library system as a user. In doing so, an instance of the User object is created to describe who that user is and what privileges he or she has been assigned.

When users are added, they are automatically included as a permanent member of a universal group (*General Users* by default). You may also assign the user to the administrative group (*Administrators* by default). All users working as members of the Administrators group have the highest level of access rights and system-wide privileges. The General Users group and the Administrators group are created automatically when you install a library system.

Note During installation you can assign custom names to the universal and administrative groups in each library system at your site. However, once you have installed the library system, you cannot rename or delete either of these groups.

Using CS Admin Tools or Web Admin Tools, you can also create two groups, zUserGroupAdmin and zPasswordAdmin, whose member users have limited administrative rights. Using Web Admin Tools (not CS Admin Tools), the users in zUserGroupAdmin can add, delete, and modify CS users and groups; the users in the zPasswordAdmin group can modify passwords of CS users and groups. However, no user in either group will be able to do any of the following:

- Delete or modify the initial Admin user or other “special” users in CS—specifically, the DirSync, ReplDaemon, and WorkflowDaemon users.
- Delete, modify, or add members to the initial Administrators group, the zUserGroupAdmin, or the zPasswordAdmin group.

If you have added other groups to the library system, you can assign a user as a member of one or more of those groups. Every group that you add to the library system has an instance of the Group object that describes the group, identifies the group members, and indicates the assigned privileges of the members. Groups are typically made up of users who work closely together and have common access requirements to documents in the library system. For example, for a library system oriented towards word-processing documents, you might create DataEntry, Supervisors, and Managers groups.

Refer to “[Library System Security](#)” on page 33 for a discussion of access control.

Document Storage and Archiving

Most users of a FileNet CS system will probably have one basic objective in mind: to manage documents efficiently. To meet this objective, they must first add their documents to a FileNet CS system. They can forget about networks, directories, path locations, and file names. One of the greatest advantages of a FileNet CS system is that users don't have to know anything about file storage to get their work done.

As an administrator, it's your job to determine the best way to distribute document storage across the network.

Storage Repositories and Storage Categories

Storage repositories are the directories or partitions you create for storing documents added to a library system. When a storage repository is added to a library system, a corresponding Storage Repository object instance is created. It contains several descriptive properties, including the system-assigned storage repository ID, the directory path, and the network name of the server on which the storage repository resides.

Storage categories are groupings of documents that have common storage requirements. You define the storage categories that best suit a particular library system. For example, if you are administering a library system largely oriented to word-processing, you could define its storage categories as Documents, Graphics, and Memos. On the other hand, if you are administering a library system oriented to financial data, you could define its storage categories as Spreadsheets, Charts, Schedules, and Miscellaneous.

Your job is to link the storage categories to the appropriate storage space, as represented by the storage repositories you assign. For example, if you were administering the word-processing library system mentioned above, you could assign the Documents and Graphics storage categories to a storage repository on a device with a large amount of space, and assign the Memos storage category to a storage repository on a device with less space allocated.

Archive Repositories and Archive Categories

Since each document in a library system has one or more versions, the amount of data in a system's storage repositories could eventually grow beyond reasonable limits. Each library system uses parameters you establish to begin the process of moving offline some (or all) document versions.

Archive repositories are the directories or partitions you create to temporarily store the versions that are moved offline by the archiving process. You can manually archive versions, or you can specify an online version limit and let the library system handle the process for you automatically.

Like storage categories, archive categories determine where a library system moves the archived document versions. For each library system, you can assign the archive categories unique names appropriate to the offline storage capabilities of that library system. For example, for a library system that archives files onto tape, sends them to a printer, or simply deletes them, you might set up archive categories named: Tape Archive, Print Archive, and Delete Without Backup. For another library system you could assign names such as Six-Month Hold, Old Docs, and Trash.

Whether or not a user can reclaim (to online status) the document versions handled by these archive categories depends on the needs or limitations of the particular library system. You can designate each archive repository as reclaimable or nonreclaimable. A library system also provides an option for the archive process to immediately delete versions from specified nonreclaimable archive repositories.

Searching for Information

Just as users don't need to know anything about document storage to add documents to the library system, neither do they need to know where the library system stored their document to be able to access them. Instead of forcing users to access information based on path locations and file names, users find documents simply by describing them.

For example, suppose a user added a document yesterday and gave it these property values:

| Property | Value |
|------------|-------------------|
| Author | Jean Giono |
| Title | John Berger and I |
| Department | Pyrenees |

Now any user with the appropriate access rights can find the document by searching on any of the values listed above. For example, you could search for all documents whose author is Jean Giono. This search would return a list of all the documents he had ever added to the library system. By adding to your search criteria, you can narrow this search to request only those documents added to the Pyrenees department by Jean Giono. The more specific your search criteria, the more precise your search results will be.

Provided you have the appropriate access rights, you can check out a version of the document listed in your search results. The library system delivers a working copy of the version to your client workstation.

Most, if not all, library system client applications provide searching capabilities based on properties. They may also provide search facilities based on the content of documents.

Content Searches

With the Content Search Manager installed in your library system, users can search for files based on words or phrases in the file contents. Users can also search XML documents for specific tags and content within tags.

The Content Search Manager supports the following content search features:


| Feature | Description |
|-------------------------|---|
| Basic Boolean searching | Allows use of AND, OR, and NOT search operators to indicate relationships between query statements. |
| Proximity searching | Allows queries for text found within a specified number of words. |
| Wildcard searching | Allows use of * and ? as wildcard placeholders in queries for all words or phrases containing specified portions of text. |
| Stopword list | Provides a list of words that will not be included in any searches, such as <i>a</i> , <i>an</i> , and <i>the</i> . |



| Feature | Description |
|------------------------|---|
| File filters | Allow files created in particular applications (for example, Microsoft Word) to be indexed in ASCII format for content searches. The third party filters which the library system uses during indexing have a limited ability to index ASCII files. Some ASCII files may not be recognized and will be passed over. |
| Embedded variants file | Allows searching for word variants based on character substitution. |
| XML searching | Allows searching content for full text, specific XML tags (for example, <title>), and content within tags. |

Version Tracking

Each time you add a document to a library system, you create the first version of that document. Each time you check out and check in the document, you create another version. The following figure illustrates this process.



Whenever you check out a document, the library system ensures that you are always working with the latest version of the document (unless you specify otherwise).

The version history of your documents is protected, because the library system delivers a copy of the document version you check out. This is important in situations where you may accidentally delete an important section of a document. If this happens, you can always cancel the first checkout, check out a new copy of the file, and make the necessary repairs. With this safety feature, the library system ensures that the original document for each version of a protected item is safely stored on the network.

As an administrator, you can also cancel the checkouts of other users. For example, you might find this feature useful if a user has left the company and forgotten to check in all his or her checked out versions of particular items.

Version Lines

Each time you use IDM Desktop (or some other library system client application) to check out the current version of a document and then check it back in, you are adding a new version of the document. The sequence of versions created by this process, as shown in the diagram above, is referred to as the primary version line. In the primary version line, versions of a document are sequentially numbered in the order in which they were checked in (1 through 9999). By default, you receive the current version when you check out a document.

Note Secondary version lines are not supported by FileNet client applications, but may be supported by customer-developed client interfaces. Refer to the *CS Explorer Online Help* for information about the *Allow Secondary Version Lines?* property of the Item object.

Library System Security

Another powerful feature of the library system is the security it provides for your important information. Each time you or any other user adds a document to the library system, you determine the document's descriptive properties and its associated version files and properties that users and groups will be able to access. You provide this information using the document's Access List property. For further security, there is also an Access List property for all users and groups.

Three objects commonly referenced by users and administrators contain an access list: the User object, the Group object, and the Item object. However, access rights to entities without access lists (such as versions) are passed down through "parent" objects (such as the Item object). So, the library system provides a simple way to set up and maintain access control, yet allows this security to be as controlled or permissive as necessary.

There are five levels of access rights in a library system:

| Access Right | Privileges Granted |
|--------------|--|
| None | None. If no other access level is stated in the access list, access rights of None are assumed. An access level of None can also be explicitly stated in an access list. |
| Viewer | Generally, the ability to view the object properties or to make copies of the associated versions. |
| Author | (Applies to documents and versions only) Viewer access rights <i>plus</i> the ability to check out, check in, and copy associated versions and modify property values for the version. In addition, you may be allowed to modify designated custom property values for the document. |
| Owner | Author access rights <i>plus</i> the ability to delete documents, modify security, and modify most properties. |
| Admin | Owner access rights <i>plus</i> the ability to modify all property values. Active members of the Administrators group are automatically assigned Admin access rights to all properties, even though their names do not appear in any access lists. Users who are not members of the Administrators group can be explicitly assigned an Admin access level to the properties associated with particular objects. |

Controlling Access through Groups

Although an access list can control access on a user-by-user basis, the lists would become long and difficult to maintain if they listed all possible users by name. To make access control easy to implement and maintain, the library system allows you to assign access rights to groups, because groups contain sets of users who have common access requirements.

Because a user may belong to more than one group in a library system, access control is determined by a user's *active group*, that is, the group in which a user is currently working. Thus, when users change their active group, their access rights to objects also change accordingly.

Using groups to control access greatly simplifies security. Suppose, for example, that there are seventy users in a library system, but these users all fall into three work groups: employees, supervisors, and a personnel department. Since there are also three groups with those same names, the access list for a document that corresponds to an employee's personnel record would need only three entries (one for each group), rather than seventy (one for each individual user).

If groups are created to match the way people work in an organization, their use in controlling access can ensure people will read and modify only the appropriate information.

How Access Lists Work

An access list contains information about the users and groups who have been assigned specific access rights to a particular object.

In CS Explorer, this information is displayed as Access Control object properties. Each Access Control object contains three pieces of information:

- The name of a user or a group
- The type of name
- The access rights granted

Here are several Access Control objects:

| Name | Type | Access Level |
|------------------|-------|--------------|
| General Users | Group | Viewer |
| Tech Pubs | Group | Author |
| Leads | Group | Owner |
| Admin Assistants | Group | Viewer |
| Sally Kemp | User | Owner |

Note Other interfaces may represent this information differently, but the access information stored in a library system is the same.

Notice the group called General Users in the example. Every user that is added to a library system is automatically made a member of the General Users group. This group provides a mechanism for adding all users to an access list in one simple entry. Another unique group in the library system is the Administrators group, whose active members have Admin access rights to all objects, even if the group is not specifically included in access lists.

Note During the installation of a library system, you can assign your own customized name for the General Users group and for the Administrators group in each library system.

By default, a library system uses the following search pattern when determining a user's access rights to a particular object and assigns the first access rights that apply:

1. Checks to see if the user is a member of the Administrators group.
2. Checks in the object's access list for an entry under the user name.
3. Checks in the object's access list for an entry under the user's active group.
4. Checks the access rights assigned to the General Users group, if included in the access list.
5. Assigns the user access rights of None.

Notice that this search pattern ensures that active members of the Administrators group always receive Admin access; thereafter, access rights specifically granted to a user take precedence over those granted to a group.

Access List Defaults

To ensure that each access list contains initial entries, a library system provides some defaults in the User, Group, and Item object. In User and Group object access lists, these defaults are standard entries that are always added, as shown in the following:

| Name | Type | Access Level |
|-----------------|-------|--------------|
| (Added By User) | User | Admin |
| (User's Name) | User | Owner |
| General Users | Group | Viewer |

The **Added By User** value is the user who added the object to the library system. This user and any user with Owner or Admin access rights can modify the values after the User or Group object has been created. The Administrators group always has Admin access rights, even though these rights are not displayed in the access list.

Default Item Access Lists

Another way a library system can help users control access to their files is by inserting default entries in document access lists so that the user does not have to provide the same set of entries each time he or she adds a document.

You can specify a set of default item access list entries for each user in his or her User object. Then, each time that user adds a document to the library system from any user interface, the access list of the Item object is filled in with those specified defaults. Of course, to cover special cases, users can always change the access list of any documents they add, but they do not have to start from scratch with each document.

In the same way that you add entries to the default item access list in the User object, you can also specify them in Group objects and the System object.

Thus, the access lists of documents do not necessarily have a standard set of default entries. The library system does add default entries to the document's access list as the document is added, but to determine these defaults, the library system will check for entries in the Default Item Access List properties in the following objects and use the first such list with any entries:

1. The User object
2. The Group object for the user's active group
3. The System object

To understand how default item access lists work, consider the following example. At the law firm of Hunter and Bowers, senior lawyer Sarah Black is using IDM Desktop to add an item to a library system. She knows that her system administrator has given her active group (Attorneys) the default item access list (shown below), which will be applied to all documents added by the group's members.

| Name | Type | Access Level |
|------------|-------|--------------|
| Attorneys | Group | Author |
| Managers | Group | Author |
| Paralegals | Group | Viewer |

The system administrator has left the Default Item Access List property blank in Sarah's User object because Sarah always wants the default access list for her active group to be applied. (Likewise, whenever Sarah changes her active group, the library system will apply the default item access list of her new active group.) She also realizes that since the default access list at her active group level does not mention her name specifically, she will receive the program default access rights for any user who adds a document and Owner access rights to the document that she is adding to the library system. And with Owner access rights to the new document, she can modify the entries in the document's access list at a later time if necessary.

Special Privileges for Custom Properties

You can assign access privileges for users to the custom properties in an Item, Version, or Checkout object (properties the administrator has defined specifically for the site) that are different from the access rights you assign them for the rest of the properties in that object. This tactic allows you to specify just what privileges (none, read, read/write, or write) users with the same access rights have to each custom property.

For example, consider a small chemical company where the employees fall into three groups: Management, Technicians, and Clerical. At this firm, a protected item is added to the library system for each chemical formula produced. The Item objects for these formulas contain several custom properties, but each group needs to read or fill out only a few values. Assigning special privileges to these properties ensures that the right group fills out the right values and that information can remain confidential.

The assigned privileges are:

| Custom Property | Group | Privileges |
|-----------------|--------|---------------------|
| Formula Name | Viewer | See (Read) |
| | Author | Modify (Read/Write) |
| | Owner | See (Read) |
| Toxicity Level | Viewer | Not see (None) |
| | Author | Modify (Read/Write) |
| | Owner | See (Read) |
| Status Assigned | Viewer | Not see (None) |
| | Author | Not see (None) |

| Custom Property | Group | Privileges |
|------------------------|--------|---------------------|
| Enforced Security Code | Owner | Modify (Read/Write) |
| | Viewer | Not see (None) |
| | Author | Not see (None) |
| | Owner | Modify (Read/Write) |

Note Refer to “Adding Custom Properties” on page 259 for suggestions and recommendations when adding custom properties.

Secure Document Delete

With the Secure Document Delete feature you can *scrub* your sensitive document files to prevent any possibility of recovery via disk recovery tools. Scrubbing a file means to overwrite its contents with a byte pattern before actually deleting them.

Refer to “Secure Document Delete” on page 331 for details about functionality.

Refer to “Storage Manager Tab” on page 167 for configuration procedures.

Clustering

A *cluster* is two or more servers (nodes) that run a common set of applications and appear as a single (virtual) server to client applications. The nodes are physically connected to one or more shared hard drives and to both a private and public network. The private network is dedicated to inter-node communication. If one of the nodes or services fails, then the cluster software enables the other node to pick up the workload of the failed node, a clustering technique called *failover*. This is the only type of clustering that Content Services supports.

When Content Services is installed in a cluster configuration, failover support makes the nodes appear to client applications as though a single *virtual server* is running Content Services. So the CS client applications see only the single Content Services virtual server and are unaware of which node is actually processing a database request.

In a cluster configuration, Content Services operates in an *active/passive* configuration. This means there is one copy of the Content Services software installed on a disk drive shared by all nodes in the cluster.

The cluster controls the operation of resources such as disk drives and applications. A resource group is a collection of logically related resources that are managed by the cluster and resident on a single node. For example, a Content Search Manager and its associated Storage Manager might compose a single *resource group*. A resource group is owned by only one node at a time.

Microsoft Cluster Server (MSCS) is normally used to administer a Windows-based cluster, but MSCS has no specific knowledge of Oracle databases and applications. So the preferred way to administer a cluster containing Oracle software is to use Oracle Fail Safe Manager (OFS), which is may be installed either within or outside the cluster being administered.

When a cluster service operation, such as failover, is performed on a resource group, the operation affects all the resources in the group. If an application within a resource group fails, MSCS can transfer the entire resource group to another node in the cluster. The failover can occur automatically (the result of an unexpected hardware or application error), or it can be initiated by whoever administers the cluster.

When a node in a cluster returns to normal operation, a *failback* operation may be executed to return a resource group back to the recovered node.

Planning Guide

For information on:

How to use this guide to plan a library system

Guidelines about the types of information you will need and who to ask

Guidelines about developing an implementation plan

A set of forms you can print and fill out when developing your implementation plan

Refer to:

“About the Planning Guide” on page 40

“Gathering Information” on page 41

“Designing Your Library System” on page 42

“Configuration Worksheets” on page 43

About the Planning Guide

NOTE Throughout this guide are references to worksheets that are provided to help you gather and analyze data, and then to plan and implement your FileNet Content Services (CS) system. Refer to “[Configuration Worksheets](#)” on page 43 for the actual forms you can print.

Who Should Use this Guide

This guide is intended for Content Services system administrators and network planning personnel. It assumes you have:

- A basic understanding of networks.
- An understanding of the planning and configuration issues associated with networks.
- A solid working knowledge of the FileNet products you will be using.

We strongly recommend attending the appropriate FileNet Training classes.

Purpose of this Guide

This guide helps you determine the information storage and access needs of your corporation and create a design for implementing a system.

To do this, you will need to examine how information flows within your organization and what corporate information model best describes your company. While this is not necessarily a complex task, it must be done before you can properly design a system that meets your needs.

How to Use this Guide

This guide describes how to use the series of “[Configuration Worksheets](#)” on page 43 that lead you through the planning process for your FileNet CS system.

Throughout this guide, the word *document* will be used to describe any information that can be stored as an electronic file (word processing, spreadsheet, graphic, image, sound, project).

If you still have questions or concerns after reading this guide, please contact your FileNet representative.

Gathering Information

The first step in planning your Content Services (CS) installation is to gather information about the documents in your organization and how they move among users. The most effective way to do this is by interview, using the provided Configuration Worksheets. While this may seem time-consuming, we have found that it is the best way to consolidate and organize the details necessary for a successful library system implementation.

The goal here is to gather information about the nature and flow of information in your organization. This will give you the resources to design a library system that will meet your specific needs. Sites that have not taken the time to gather this information before beginning to implement have generally had to modify their systems after their initial setup.

Some of the information you will need is general descriptive data, intended to help you make decisions about your overall library system configuration (such as budgetary constraints that force you to design your system using one Property Manager for your entire site). Other information will help you to set up specific property values within your library system (such as the names of your users, the groups to which they belong, and what access rights they will need to specific types of documents).

Whom to Interview

You may want to consider interviewing some of the following sources:

- Division Heads
- Department Heads
- Administrative Assistants
- Word Processing Managers
- Delivery and Document Handling Personnel
- Telephone System Manager
- Network Administrators (WAN and LAN)
- Receptionist

What to Ask

The Configuration Worksheets are designed to drive the interview process. Refer to “[Configuration Worksheets](#)” on page 43 for guidelines on using these forms.

Designing Your Library System

The following topics, in conjunction with the “[Configuration Worksheets](#)” on page 43, lead you through the steps for designing your FileNet Content Services (CS) system. You will use the general and specific information gathered in the previous topic to assist you with your design process. Once you have completed these worksheets, you will find that the difficult part of the installation process--a good plan--is done.

Describe and Diagram Your LAN

The first step in actually designing your FileNet CS system is to use the “[LAN Description Worksheet](#)” on page 57 to describe and diagram the major components of your LAN as they are planned for the proposed system installation. FileNet CS system configuration and performance depend on the network as well as the client and server hardware and software configuration.

As you create your diagram, ask yourself the following questions:

- What do I have to do, if anything, to my current network to support a FileNet CS system at the desired performance level?
- Do I need to add or upgrade any servers (including web servers)?
- Do I need to add or upgrade network links (such as communication lines, bridges, routers, hubs, segments, rings, etc.)?
- Do I need to add or upgrade any clients?
- Will my FileNet CS system be based on one library system? Or do I want to share the load among multiple library systems, making sure that all users have access to the information required to do their work?

Estimate Hardware and Software Requirements

In the “Hardware and Software Requirements” worksheets, you will provide the detail necessary to actually install and configure your FileNet CS system. Use the information you gathered in the previous topic, “[Gathering Information](#)” on page 41.

After you have completed these configuration worksheets, draw a second diagram of your LAN, this time placing your FileNet CS system within it. Check the loading and access to be sure they can meet your access and performance expectations.

Configuration Worksheets

The configuration worksheets are provided to help you gather and analyze the data you will need to plan your library system. As you interview your sources, you should ask the following questions.

Who?

Use the “Answering Who” worksheets when asking:

- Who will access the library system? Who does each user work with? How do these people work together (departments, project teams, functional or cross-functional work groups)? Refer to “Users and Groups Worksheets” on page 45.
- Who will act as installer? System manager? Site administrator? Sub-administrator? Refer to “Administrators Worksheet” on page 47.

What?

Use the “Answering What” worksheets when asking:

- What are the types of documents and what are their requirements? Refer to “Document Description Worksheet” on page 48.
- What are the access rights of specific users/groups to each type of document? Refer to “Document Access Rights Worksheet” on page 49.
- What is the network setup? What types of WANs, LANs, servers, and client machines are used? (This information is necessary because a CS system depends upon the performance and integrity of the network to function properly.) Refer to “Network Worksheet” on page 51 and “LAN Description Worksheet” on page 57.
- What are your hardware and software requirements? Refer to:
 - “Machine Requirements Worksheet” on page 58
 - “Library System Components and Sizing Worksheet” on page 59
 - “Client Application Requirements Worksheet” on page 61
 - “Library System Requirements Worksheet” on page 62
 - “Initial Storage Server Hardware Worksheet” on page 65
 - “Additional Storage Server Hardware Worksheet” on page 66
- What are the budgetary constraints? Refer to “Budgetary Constraints Worksheet” on page 52.

Where?

Use the “Answering Where” worksheet when asking:

- Where are the documents stored and accessed?
- Do all these locations have connections between them? What is the availability of these connections?
- Is remote access required?
- Are there other special considerations?

Refer to the “[Document Storage & Access Worksheet](#)” on page 53.

When?

Use the “Answering When” worksheets when asking:

- When is each site in your organization open for operation? Refer to “[Site Hours Worksheet](#)” on page 54.
- When is your network used most? least? Refer to “[Network Usage Worksheet](#)” on page 55.
- When are documents accessed at a given location? Refer to “[Document Timeframes Worksheet](#)” on page 56.
- When a disaster occurs, what is your recovery plan? Refer to “[Disaster Recovery Worksheet](#)” on page 67.

Users and Groups Worksheets

Use this table to help you determine what users and groups to add to your FileNet CS system. List all the groups in your organization across the top of the table. Then list each user name in the first column and place a check mark under each group to which the user belongs.

| User Names | G r o u p | | | | | | | | | | | | | | | |
|------------|-----------------------|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |

Total # of users in organization:

Use this table to indicate how many users belong to each of the groups in your organization.

| Group Name | # of Users |
|-------------------|-------------------|
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |

Document Description Worksheet

Use this table to determine the categories of documents in your organization and their specific requirements. List the document categories on the top (such as financial spreadsheets, engineering specs, manuals) and then provide the necessary information for each document category.

| | D o c C a t | | | | | | | | | | | | | | |
|---|--------------------------------|--|--|--|--|--|--|--|--|--|--|--|--|--|--|
| Mission Critical, General Office, or Reference? | | | | | | | | | | | | | | | |
| Total # of files | | | | | | | | | | | | | | | |
| Predicted # of files over next year | | | | | | | | | | | | | | | |
| Predicted # of versions per file | | | | | | | | | | | | | | | |
| Avg size of file | | | | | | | | | | | | | | | |
| Storage Requirements | | | | | | | | | | | | | | | |
| # of unprotected items | | | | | | | | | | | | | | | |
| Performance expectations/needs | | | | | | | | | | | | | | | |

Document Access Rights Worksheet

Use this table to determine default item access lists for each document category. List each document category across the top. In the first column, list the groups and, when necessary, specific users who require access to these documents. Then indicate the security access rights of each group and user to the document category.

| Group / User Names | D o c u m e n t C a t e g o r y | | | | | | | | | | | | | | | |
|-----------------------|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|--|
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |
| | | | | | | | | | | | | | | | | |

Access Rights

| Level | Rights Granted |
|--------|--|
| Admin | Owner access rights <i>plus</i> the ability to modify all property values. |
| Owner | Author access rights <i>plus</i> the ability to modify most property values, including security, and to delete the document. |
| Author | Viewer access rights <i>plus</i> the ability to check out and check in the document and change a few property values. |
| Viewer | The ability to view and make copies of the document and view the properties. |
| None | None. The user doesn't even know the object exists. |

Network Worksheet

Describe the following aspects of your existing network setup.

| | |
|---|--|
| Network topology | |
| | |
| WAN | |
| Speed | |
| Performance | |
| Availability | |
| Bridges & routers | |
| Remote access - freq. & quantity of use | |
| | |
| LAN | |
| Network OS type & version | |
| Server speed | |
| Performance | |
| Network interface cards | |
| Server IP addresses | |
| Disk space | |
| Backup | |
| Freq. & quantity of use | |
| | |
| Client Setup | |
| Hardware | |
| Configurations | |
| Run software across network? from local disk? | |
| Web browser | |

Budgetary Constraints Worksheet

List any budgetary constraints that could affect your library system setup.

| | |
|-----|--|
| 1. | |
| 2. | |
| 3. | |
| 4. | |
| 5. | |
| 6. | |
| 7. | |
| 8. | |
| 9. | |
| 10. | |

Document Storage & Access Worksheet

List the document categories on the top (such as financial spreadsheets, engineering specs, manuals) and then provide the necessary information for each document category.

| | D o c C a t | | | | | | |
|--|--------------------------------|--|--|--|--|--|--|
| Storage requirements | | | | | | | |
| Storage location | | | | | | | |
| Access location | | | | | | | |
| Connections between storage and access sites | | | | | | | |
| Availability of connections | | | | | | | |
| Remote access? | | | | | | | |
| Special considerations | | | | | | | |
| Web access | | | | | | | |

Site Hours Worksheet

Indicate the typical work hours of the site where the library system will be installed.

| Sites | Hours of Operation |
|-------|--------------------|
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |
| | |

Network Usage Worksheet

Indicate the usage patterns of your network.

| |
|--|
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |
| |

Document Timeframes Worksheet

List the document categories in the first column (such as financial spreadsheets, engineering specs, manuals) and then provide the necessary access information for each document category.

| Document Category | Time of file access | Freq. of file access | Avg. life of files |
|-------------------|---------------------|----------------------|--------------------|
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |
| | | | |

LAN Description Worksheet

In addition to completing the following LAN description worksheet, make sure you have a complete LAN diagram.

| User/Network Info | Number | Descriptions |
|----------------------------|--------|--------------|
| User Information | | |
| # of LAN users | | |
| Types of users | | |
| Types of files | | |
| # of files | | |
| Avg. size of files | | |
| Freq. of access | | |
| Current response time | | |
| | | |
| Network Information | | |
| NW OS type & version | | |
| Total # of nodes | | |

Machine Requirements Worksheet

Indicate the number of computers required and a brief description of each type.

| Component | Quantity | Descriptions |
|---|----------|--------------|
| Initial storage servers | | |
| Additional storage servers | | |
| Verity K2 Master Admin servers | | |
| Standalone database servers (optional, if supported) | | |
| Content Services clients | | |

Library System Components and Sizing Worksheet

System #1 Name:

| Specific Component | Value | Comments |
|---|-------|----------|
| Database server (name) | | |
| Standalone? (yes or no) | | |
| Size of database or tablespace | | |
| Size of database transaction or redo log | | |
| Rollback database segment tablespace (Oracle) | | |
| Temporary database tablespace (Oracle) | | |
| Physical location | | |
| Archive media | | |
| # of users (w/ growth) | | |
| # of storage repositories | | |
| Disk space available | | |
| Physical location | | |
| Projected # of files (w/ growth) | | |
| Initial storage server (name) | | |
| FileNet CS system name | | |
| # of storage repositories | | |
| Disk space available | | |
| Physical location | | |
| Projected # of files (w/ growth) | | |
| # of content search repositories | | |
| Projected content search repository sizes | | |

| Specific Component | Value | Comments |
|---|-------|----------|
| Additional storage server (name) | | |
| # of storage repositories | | |
| Disk space available | | |
| Physical location | | |
| Projected # of files (w/ growth) | | |
| # of content search repositories | | |
| Projected content search repository sizes | | |

Client Application Requirements Worksheet

| Client Application | Version | Configuration Requirements |
|--------------------|---------|----------------------------|
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |
| | | |

Library System Requirements Worksheet

List the overall system requirements for your library system

| System Properties | Value | Comments |
|------------------------------|-------|----------|
| Online version limit | | |
| Permanent versions in limit? | | |
| Secure searches? | | |
| Default Checkout File Name | | |
| Custom property 1 | | |
| Custom property 2 | | |
| Custom property 3 | | |
| Custom property 4 | | |
| Custom property 5 | | |
| Custom property 6 | | |
| Custom property 7 | | |
| Custom property 8 | | |
| Custom property 9 | | |
| Custom property 10 | | |
| Storage categories | | |
| 1. | | |
| 2. | | |
| 3. | | |
| 4. | | |
| 5. | | |
| 6. | | |
| 7. | | |
| 8. | | |
| Archive categories | | |
| 1. | | |
| 2. | | |
| 3. | | |

| System Properties | Value | Comments |
|-------------------|-------|----------|
| 4. | | |
| 5. | | |
| 6. | | |
| 7. | | |
| 8. | | |

Database Server Hardware Worksheet

| Property Server | Model/Ver. | Description | Size/Speed |
|------------------------------|------------|-------------|------------|
| Server brand & model | | | |
| CPU (speed & number) | | | |
| Memory | | | |
| Storage | | | |
| Network interface card | | | |
| Archive media (size & speed) | | | |
| OS version | | | |

Initial Storage Server Hardware Worksheet

| Property Server | Model/Ver. | Description | Size/Speed |
|------------------------------|------------|-------------|------------|
| Server brand & model | | | |
| CPU (speed & number) | | | |
| Memory | | | |
| Storage | | | |
| Network interface card | | | |
| Archive media (size & speed) | | | |
| OS version | | | |

Additional Storage Server Hardware Worksheet

| Server | Model/Ver. | Description | Size/Speed |
|------------------------------|------------|-------------|------------|
| Server brand & model | | | |
| CPU (speed & number) | | | |
| Memory | | | |
| Storage | | | |
| Network interface card | | | |
| Archive media (size & speed) | | | |
| OS version | | | |

Disaster Recovery Worksheet

| Disaster Recovery | Value | Comments |
|--|-------|----------|
| Disk duplexing / mirroring / disaster recovery | | |
| Off-site backup storage | | |

Roadmap to Setting Up

For information on:

Installing and configuring the prerequisite software, installing the initial library system services, and adding additional services or a display language to an existing library system

Installing and configuring the administrative client

Configuring a library system to meet the specific requirements of your site

Installing client software for accessing a library system

Upgrading the database and Content Services software

Refer to:

“Installation Roadmap” on page 69

“Administrative Client Setup Roadmap” on page 71

“Configuration Roadmap” on page 72

“Client Interface Setup Roadmap” on page 74

“Upgrade Roadmap” on page 75

Installation Roadmap

To install a library system:

| | |
|--------------------|---|
| Required | <p>Before you begin any library system installation or upgrade, always check the FileNet Worldwide Customer Support web site for documentation updates and software patches. See “About the Documentation” on page 12 if you need help.</p> <p>If you plan to install IDM Desktop and CS 5.5 server components on the same Windows machine, install IDM Desktop first.</p> <p>Note To configure and administer a library system, you must set up an administrative client. Refer to “Administrative Client Setup Roadmap” on page 71.</p> |
| Highly Recommended | <p>For an introduction to Content Services architecture and library system functionality, read these topics in the <i>Overview</i>:</p> <ul style="list-style-type: none"> • “What is Enterprise Content Management?” on page 19 • “The Library System” on page 22 <p>To familiarize yourself with some of the main features, read these overview topics in the <i>Administration Guide</i>:</p> <ul style="list-style-type: none"> • “Storage Repositories” on page 233 and “Storage Categories” on page 240 • “Archive Repositories” on page 242 and “Archive Categories” on page 250 • “Content Search Repositories” on page 253 • “Custom Properties” on page 258 • “User Objects” on page 263 • “Group Objects” on page 267 • “Adding a Display Language” on page 134 |
| Highly Recommended | <p>Before beginning installation of the initial library system services, complete the worksheets included with the “Planning Guide” on page 39. This will help you plan your library system in advance, and will provide a detailed record of your library system configuration.</p> |
| Highly Recommended | <p>Complete the “Installation Settings” on page 119. This will create a record of your installation settings before you start installing Content Services.</p> |
| Required | <p>Review the “Release Notes” on page 17 for a list of new features, known problems, and other critical information pertaining to this release.</p> |

- Required Verify that all platform (hardware and software) requirements have been met for the server on which you plan to install the library system. Install any necessary updates or service packs. See [CS 5.5.0 Windows Oracle Platform Requirements at FileNet Worldwide Customer Support](#).
- Required Install and configure the database engine and prepare the database, as specified in [“Setting Up the Oracle Software” on page 89](#)
- Required Install the Content Services software, as specified in [“Installing Initial Library System Services” on page 123](#).

To install additional library system services:

- Highly Recommended Familiarize yourself with the general installation and configuration issues of the library system services. In addition, identify the number of storage, archive, and content search repositories to configure on the additional storage server as well as the storage and archive categories that link to these repositories.
- Highly Recommended Complete the [“Installation Settings” on page 119](#). This will create a record of your installation settings.
- Required Verify that all platform (hardware and software) requirements have been met for the server on which you plan to install additional services. Install any necessary updates or service packs. See [CS 5.5.0 Windows Oracle Platform Requirements at FileNet Worldwide Customer Support](#).
- Required Install the additional library system services as specified in [“Installing Additional Library System Services” on page 130](#).

To install Rendition Services:

- Highly Recommended Familiarize yourself with the general requirements and installation procedures as described in these topics:
 - [“Hardware and Software Requirements for Rendition Services” on page 659](#)
 - [“Install Rendition Services on Servers” on page 664](#)

To configure Replication Services:

- Highly Recommended Read the following topics and determine if an administration domain is required before configuring replication:
 - [“Replication Services Overview” on page 594](#)
 - [“Planning for Replication Services” on page 601](#)
 - [“Replicating Objects” on page 609](#)

Note To configure and administer a library system, you must set up an administrative client. Refer to [“Administrative Client Setup Roadmap” on page 71](#).

Administrative Client Setup Roadmap

The administrative client software for Content Services is CS Admin Tools. Also available is a web browser-oriented variation called CS Web Admin Tools, which provides a lighter set of administrative tools, and also the ability to start/stop CS services remotely via a web browser.

To set up your administrative client:

| | |
|----------|--|
| Required | Verify that all platform (hardware and software) requirements have been met for the client machine on which you plan to install the administrative client. Install any necessary updates or service packs. See CS 5.5.0 Windows Oracle Platform Requirements at FileNet Worldwide Customer Support . |
| Required | Install IDM Desktop on the administrative client machine. |
| Required | Install Content Services Admin Tools on the administrative client. |
| Optional | On a web server, install Content Services Web Admin Tools and CS Client Library Files (the latter is installed by FileNet Web Services or IDM Desktop). |

Note Before you begin using your library system, there are a number of configuration issues to consider. Refer to “[Configuration Roadmap](#)” on page 72.

Configuration Roadmap

Note For information about performing administrative tasks for your library system, refer to “Administration Guide” on page 222.

To configure your library system:

| | |
|--|--|
| Highly recommended | Familiarize yourself with important configuration issues. Refer to “Configuration And Maintenance” on page 224. |
| Highly recommended | Create additional repositories for storing your documents. Refer to “Storage Repositories” on page 233 and “Adding a Storage Repository” on page 234. |
| Required if you add storage repositories | Create additional groupings of documents that have common storage requirements. Refer to “Storage Categories” on page 240 and “Adding a Storage Category” on page 241. |
| Highly recommended | Create additional repositories for archiving your documents. Refer to “Archive Repositories” on page 242 and “Adding an Archive Repository” on page 247. |
| Required if you add archive repositories | Create additional groupings of documents that have common archive requirements. Refer to “Archive Categories” on page 250 and “Adding an Archive Category” on page 251. |
| Required for Content Search Managers | Create additional repositories for the index files used in content-based searches. Refer to “Content Search Repositories” on page 253 and “Adding a Content Search Repository” on page 255. |
| Required if you use Document Classes | Add custom properties to document, version, and checkout object types. Refer to “Custom Properties” on page 258 and “Adding Custom Properties” on page 259. |
| Optional | Add users. Refer to “User Objects” on page 263 and “Working with User Objects” on page 263. Note To bulk add users, refer to <i>CS Bulk Process Online Help</i> . Also see “Importing Users and Groups” on page 272 for information about importing users from your network domain or LDAP directory to your library system. |
| Optional | Add groups. Refer to “Group Objects” on page 267 and “Working with Group Objects” on page 268. Note To bulk add groups, refer to <i>CS Bulk Process Online Help</i> . Also see “Importing Users and Groups” on page 272 for information about importing groups from your network domain or LDAP directory to your library system. |
| Optional | Edit the configuration parameters. Refer to “The CS Configuration Tool” on page 164. |
| Optional | Reconfigure database parameters to improve system performance. Refer to “Configuring Databases” on page 229. |

| | |
|----------|--|
| Optional | Configure port numbers at sites that use WANs and firewalls. Refer to “Configuring for WANs and Firewalls” on page 117. |
| Optional | Add Document Classes to specify the property display sets for associated documents. Refer to <i>CS Document Class and CVL Admin Online Help</i> . |
| Optional | Add controlled vocabulary lists (CVLs). CVLs provide required selection lists for associated custom properties. Refer to the <i>CS Document Class and CVL Admin Online Help</i> . |
| Optional | Bulk add pre-existing documents. Refer to “Managing Storage Space” on page 226 for information about avoiding disk space overflow. Refer to the <i>CS Bulk Process Online Help</i> for procedures. |
| Optional | Configure Replication Services. Refer to “Replication Services” on page 593. |
| Optional | Turn on and configure audit logging. Audit logging tracks particular library system processes. Refer to “Audit Logging” on page 323. |
| Optional | Install additional display languages. Refer to “Adding a Display Language” on page 134. |
| Optional | Set up multiple library systems at your site. Refer to “Accessing Multiple Library Systems” on page 204 and “Managing Multiple Library Systems” on page 310. |
| Required | Implement a backup strategy for each library system's database information and stored documents. Refer to “Backing Up the Library System” on page 336. |

Note For users to access a library system, they must have a client interface installed on their workstation. Refer to [“Client Interface Setup Roadmap”](#) on page 74.

Client Interface Setup Roadmap

For users to access a library system, you must install a client interface on their workstations or set up a web server through which users can access the library system using a web browser. Supported client interfaces include:

- IDM Desktop
- Web Services
- eProcess

See [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#) for a complete list of certified client interfaces. Refer to the client interface's documentation for installation instructions.

Note If your site has its own customized client interface to Content Services, refer to its documentation and to “[Client Library Files](#)” on page 207.

Upgrade Roadmap

Note If your library system is older than version 5.4, you must upgrade to the version 5.4 before you can upgrade to version 5.5.

To upgrade a library system from version 5.4 to version 5.5:

| | |
|--------------------|---|
| Required | <p>Before you begin any library system installation or upgrade, always check the FileNet Worldwide Support web site for documentation updates and software patches. See “About the Documentation” on page 12 if you need help. The documentation includes the following:</p> <ul style="list-style-type: none"> • Read CS 5.5.0 Windows Oracle Release Notes for information that may affect your upgrade. • Review all platform (hardware and software) requirements for the servers on which you plan to upgrade Content Services. See CS 5.5.0 Windows Oracle Platform Requirements at FileNet Worldwide Customer Support. |
| Required | <p>Do a full backup of the library systems you plan to upgrade. If the upgrade fails for any reason, you may need to restore the library system from this backup. Refer to “Backing Up the Library System” on page 336.</p> |
| Required | <p>Install any necessary updates or service packs indicated in the platform requirements for the servers on which you plan to upgrade Content Services.</p> |
| Required | <p>Upgrade your client interfaces.</p> <p>Note Saros Document Manager and @mezzanine are not supported for version 5.1 and higher library systems. Refer to “Client Interface Setup Roadmap” on page 74 for current client requirements.</p> |
| Required | <p>Upgrade your library system servers as specified in “Upgrading Content Services” on page 152.</p> |
| Required | <p>Upgrade your administrative client.</p> <p>Note FileShare is not supported for version 5.1 and higher library systems. The administrative client for Content Services 5.3 and higher is Content Services Admin Tools. Refer to “Administrative Client Setup Roadmap” on page 71 for current administrative client requirements.</p> |
| Highly recommended | <p>Do a full backup of the upgraded library systems. Refer to “Backing Up the Library System” on page 336.</p> |

Installation Guide

| For information on: | Refer to: |
|--|---|
| Verifying compliance with hardware, software, server name, and database name requirements, installing the database engine and creating the library system database | “Prerequisites to Installing Content Services” on page 80 |
| Installing the initial library system services | “Installing Content Services” on page 119 “Installing Content Services Unattended” on page 136 |
| Adding a Storage Manager, Content Search Manager, Replication Services, or Display Language to an existing library system | “Installing Additional Library System Services” on page 130 |
| Installing Content Services in a cluster environment | “Content Services in a Cluster Environment” on page 83 |
| Installing Client Library Files for a custom client interface, or updating or reinstalling the library files (initial installation of Client Library Files is done automatically during installation of FileNet clients and servers) | “Client Library Files” on page 207 “Installing Client Library Files Unattended” on page 208 |
| Preparing for and performing an upgrade from version 5.4 to version 5.5 on initial storage servers and additional storage servers | “Upgrading Content Services” on page 152 |
| Configuring an installed or upgraded library system to meet your site's specific needs | “The CS Configuration Tool” on page 164 “Configuring Content Services” on page 164 “Accessing Multiple Library Systems” on page 204 |
| Troubleshooting installation problems | “Problem Solving” on page 215 |
| Uninstalling library system or client software | “Uninstalling Content Services” on page 217 |

Installation Overview

A Content Services (CS) installation consists of one *database engine* (Microsoft SQL Server 2000, Microsoft SQL Server 2005, Oracle9i 9.2.0, or Oracle 10g 10.2.0.1.0) and, at a minimum, one initial set of *library services* (components), consisting of a Storage Manager, a Content Search Manager, and Replication Services. The initial set of library services and the associated CS property information stored in the database constitute a *library system* (or *library* for short).

Beyond the minimal configuration, you may choose to have the database engine support multiple library systems. To install another library system, you install another initial set of library services. You can also install additional sets of services for a given library system on other servers. Installing additional services allows you to make use of additional disk space (for library storage) and speed up processing by spreading activity over multiple CPUs.

Content Services Setup (CS Setup for short), the program that installs initial and additional library services, installs a complete *set* of services, not individual services. As noted above, an initial set of services consists of a Storage Manager, a Content Search Manager, and Replication Services. An additional set of services consists of a Storage Manager and a Content Search Manager. In other words, you cannot install just a Content Search Manager, or just a Storage Manager. So your library system will always consist of equal numbers of Storage Managers and Content Search Managers, with exactly one Replication Services component.

Some component terminology: We refer to a component, such as a Content Search Manager, as an *initial* Content Search Manager or an *additional* Content Search Manager depending on whether the component is part of a set of initial services or additional services for a library system.

Some server terminology: The server where the database engine is installed is called the *database server*. Because this server stores the properties that define a CS library system, it is also known as the *property server*. A database server where no library services are installed is called a *standalone database server*. A server where at least one set of library services is installed is called a *storage server*. The server where the initial set of services of a library system is installed is called the *initial storage server* for that library system. A server where an additional set of services of a library system is installed is called an *additional storage server* for that library system. For additional information on terminology see the “Glossary” on page 718.

Thus, a given server can simultaneously have one or more of the following labels: database server, initial storage server, and additional storage server, depending on what is installed on it. For example, a single server can acquire all three labels by hosting a database engine, the initial set of services for one library, and an additional set of services for another library whose initial services are on some other server.

Installation Constraints

CS Setup allows multiple sets of initial and/or additional library services to be installed on the same server, including the database server, provided that no two sets (initial or additional) of services are for the same library system. That is, if an initial set of services for a given library is on a server, no additional set of services for the library can be on the same server. And if one additional set of services for a library is on a server, then any other additional set of services for the same library must be on a different server.

CS Setup imposes an additional installation constraint based on the requirement that the document indexing operations of the initial and all additional Content Search Managers of a library system must be governed by a single instance of the Verity K2 Master Admin program (Master Admin for short). A Content Search Manager is said to “point to” the Master Admin that governs it. CS Setup installs (as needed) and assigns a Master Admin to govern each initial or additional Content Search Manager it installs, according to the following criteria:

- If CS Setup is installing an initial set of library services on a server where no library services are already installed, *and* CS Setup has not previously installed any library services on other servers in your organization, it automatically installs a Master Admin on this server to govern the initial Content Search Manager being installed.
- If CS Setup is installing the initial set of library services on a server where no services of any other library are already installed, *and* CS Setup has previously installed library services on other servers in your organization, it allows you to choose which Master Admin will govern the initial Content Search Manager being installed: You can have CS Setup either install a Master Admin on this server to govern the initial Content Search Manager being installed or else have the Content Search Manager point to a Master Admin already installed on another server.
- If CS Setup is installing an initial set of library services on a server where initial or additional services of other libraries have already been installed, CS Setup will automatically choose the Master Admin that governs the first (initial or additional) Content Search Manager that CS Setup installed on the server.
- If CS Setup is installing an additional set of services for a library on a server, it will try to assign the additional Content Search Manager being installed to point to the Master Admin that governs the initial Content Search Manager for the library. If some other Master Admin was chosen to govern an initial or additional Content Search Manager (for some other library) previously installed on the server, then the current CS Setup execution will fail.

The above criteria imply the constraint that all Content Search Managers residing on the same server must be governed by the same Master Admin. In particular, no two additional Content Search Managers can reside on the same server unless their respective initial Content Search Managers are governed by the same Master Admin.

Installation Scenario

To illustrate these installation constraints, suppose you run the following sequence of CS Setup executions in an organization where CS Setup has not previously run:

1. CS Setup installs the initial services for library LIB1 on server SVR1 and, without choice, installs a Master Admin on SVR1 to govern the initial Content Search Manager of LIB1.
2. CS Setup installs the initial services for library LIB2 on server SVR2, accepting your choice to install a Master Admin on SVR2 to govern the initial Content Search Manager of LIB2.
3. CS Setup installs the initial services for library LIB3 on server SVR1 and, without choice, points the initial Content Search Manager of LIB3 to the Master Admin of LIB1 on SVR1.
4. CS Setup installs the initial services for library LIB4 on server SVR3, accepting your choice to point the initial Content Search Manager of LIB4 to the Master Admin on SVR2.
5. CS Setup installs additional services for library LIB2 on server SVR4 and, without choice, points the additional Content Search Manager of LIB2 to the Master Admin of LIB2 on SVR2.
6. CS Setup installs additional services for library LIB2 on server SVR3 and, without choice, points the additional Content Search Manager of LIB2 to the Master Admin of LIB2 on SVR2.
7. CS Setup fails to install additional services for library LIB2 on server SVR1. Why? The initial Content Search Manager for LIB1, installed on SVR1 (in step 1), points to the Master Admin on SVR1. So any Content Search Manager subsequently installed on SVR1 must also point to the Master Admin on SVR1. But all additional

Content Search Managers of LIB2 must point to the Master Admin on SVR2 that governs the initial Content Search Manager of LIB2. CS Setup avoids the potential conflict by failing to carry out the installation.

8. CS Setup fails to install additional services for LIB1 on server SVR4. Why? The additional Content Search Manager for LIB2 installed in step 5 points to the Master Admin of LIB2 on SVR2. So any Content Search Manager subsequently installed on SVR4 must also point to the Master Admin on SVR2. But all additional Content Search Managers of LIB1 must point to the Master Admin that governs the initial Content Search Manager of LIB1. And that Master Admin is on SVR1. CS Setup avoids the potential conflict by failing to carry out the installation.

Installation Platforms

You can install CS 5.5.0 on any of four platforms: Windows SQL Server, HP Oracle, Solaris Oracle, and Windows Oracle. The *platform* is identified by

- The operating system (Windows, HP-UX or Solaris) running on the initial storage server and any additional storage servers
- The type of database engine (Microsoft SQL Server 2000 or Oracle)

For example, a Windows SQL platform means that Windows runs on the initial storage server and any additional storage servers, and the database engine is SQL Server. The Windows SQL platform is homogeneous in the sense that the database server and the initial/additional storage servers all run the Windows operating system.

An Oracle platform may be homogeneous, if the database server and the initial/additional storage servers all run the same type of operating system. Or it may be heterogeneous, if the database server and the initial/additional storage servers run different types of operating systems. For example, on a Windows Oracle platform, Windows runs on the initial and additional storage servers, while the database server may be running the HP-UX or Sun Solaris operating system.

Prerequisites to Installing Content Services

Before you begin the actual installation of a library system, it is important to ensure that your hardware and software resources are in order, and your Oracle database is correctly installed and configured. You should also determine what values you want to supply for the various prompts displayed during the installation.

See the “[Installation Roadmap](#)” on page 69 for a complete list of the tasks required when installing a library system or adding services to an existing library system. Before beginning an installation or upgrade, refer to:

- [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#) for the minimum hardware and software required for database servers, storage servers, and client machines, as well as a list of supported CS client applications. Check the platform requirements from time to time, as they may be updated over the period that Content Services 5.5 is supported.
- Software patches for CS 5.5.0 at the [FileNet Worldwide Customer Support](#) web site. See Software Patches under “[About the Documentation](#)” on page 12 if you need help.
- “[Windows 2000 or Windows 2003 Requirements](#)” on page 81
- “[Server Name Requirements](#)” on page 81
- “[Database Name Requirements](#)” on page 82
- “[Library System Name Requirements](#)” on page 82
- “[Setting Up the Oracle Software](#)” on page 89
- “[Creating a Database User](#)” on page 115
- “[Configuring for WANs and Firewalls](#)” on page 117
- “[Installation Settings](#)” on page 119 for a table you can use to plan and record the values you will supply when prompted during the installation of your library system.

Note At least **1 week** prior to the installation, the TSP (Technical Service Provider) **must** schedule the installation with the FileNet Upgrade/Install Assurance Team (U/I Team). To access the U/I Team's latest scheduling procedures go to [FileNet Worldwide Customer Support](#) web site.

If you are installing Content Services in a specialized environment, refer to:

- “[Installing Content Services with Image Services](#)” on page 80
- “[Installing Content Services into Non-English Environments](#)” on page 81
- “[Content Services in a Cluster Environment](#)” on page 83
- “[Installing and Configuring Microsoft Cluster Server](#)” on page 84

Installing Content Services with Image Services

Installing Content Services 5.5 to share an Oracle9i 9.2.0 database engine with an existing Image Services 3.6 installation is supported. Installation and configuration details are provided in a technical notice titled *Installing CS to Share a Database Engine with IS*.

To obtain Content Services technical notices, log on to the [FileNet Worldwide Customer Support](#) web site. See Documentation Updates under “[About the Documentation](#)” on page 12 if you need help.

Installing Content Services into Non-English Environments

Before installing Content Services 5.5 on a server with a non-English operating system or database engine, please refer to the technical notice titled, *Installing Content Services into Non-English Environments*.

To obtain Content Services technical notices, log on to the [FileNet Worldwide Customer Support](#) web site. See Documentation Updates under “[About the Documentation](#)” on page 12 if you need help.

Windows 2000 or Windows 2003 Requirements

You **cannot** install a library system on a domain controller.

In order to configure a Windows 2000 or Windows 2003 server for direct document access, users must have Read permission to the content search repository directories. To do this under Windows 2000 or Windows 2003, the Allow Read permission check box on the Share Permission tab should be checked. (The Share Permission tab is accessed by clicking the Permissions button on a directory Properties dialog box.) The check boxes for Full Control and Change should be left unchecked. If either of these are checked, users will not have access to the documents in the content search repository directory.

Content Services 5.5 provides limited Windows 2000 or Windows 2003 Active Directory support, enabling the CS Service Manager (refer to “[Controlling Library System Services](#)” on page 198) to start, stop, and monitor library system services remotely. Library system services throughout an Active Directory forest can be monitored and controlled from any library system server. This functionality requires the following configuration:

- The library system servers must be on the same IP subnet.
- The library system servers must be in the same Active Directory forest.
- The Enterprise Admins group and the Domain Admins group must belong to the local Administrators group on each library system server in the Active Directory forest.

You can find the hardware and software requirements for these client applications at [the FileNet Worldwide Customer Support](#) web site.

Server Name Requirements

Host names of storage servers and a Windows property server must not exceed 15 characters. The host name of a UNIX property server must not exceed 30 characters.

Moreover, any server where a Verity K2 Master Admin is installed (as part of the installation of the first set of initial library services on a server) must have a fully qualified domain name of no more than 255 characters. To ensure that the domain name does not exceed 255 characters, do the following before installing any initial library services on a server:

1. Right-click the My Computer icon on the Windows 2000 or Windows 2003 desktop and choose Properties.
2. On the System Properties dialog box, select the Network Identification tab and click Properties.
3. On the Identification Changes dialog box, click More.

4. On the DNS Suffix and NetBIOS Computer Name dialog box, clear the contents of the text box titled “Primary DNS suffix of this computer” and click OK.

Note Perform this step only if fully qualified domain name of the server exceeds 255 characters.

Do *not* install library system components on servers having accented characters in their computer name. The directory naming services for some operating systems change such accented characters in path names; this, in turn, adversely affects the library system's ability to locate certain system components and directories.

Database Name Requirements

Do not use *cluster* as a database name, even if you are not going to install the database engine in a cluster environment.

Library System Name Requirements

Each library system must have the same name as its tablespace. The case-insensitive tablespace name should uniquely identify this library system on the enterprise network (LAN and WAN) and must be unique from any other library system name on the enterprise network. A name can be up to 30 alphanumeric or underscore characters, but cannot start with a number or underscore.

Note Due to restrictions in Oracle tablespace naming, accented characters are *not* allowed in the tablespace/library system name.

Due to database restrictions, you cannot use the following keywords (upper or lowercase) as a library system name:

add, all, alter, and, any, arith_overflow, as, asc, at, authorization, avg, begin, between, break, browse, bulk, by, cascade, case, char_convert, check, checkpoint, close, clustered, coalesce, comment, commit, committed, compute, confirm, constrain, continue, controlrow, convert, count, create, current, current_date, current_time, current_timestamp, current_user, cursor, data_pgs, database, date, dbcc, deallocate, declare, default, delete, desc, disk, distinct, double, drop, dummy, dump, else, end, endtran, errlvl, errvl, errexit, escape, escept, except, exec, execute, exists, exit, fetch, fillfactor, floppy, for, foreign, from, goto, grant, group, having, holdlock, identity, identity_insert, identitycol, if, in, index, insensitive, insert, intersect, into, is, isolation, key, kill, level, like, lineno, load, max, min, mirror, mirrorexit, national, nocheck, noholdlock, nonclustered, not, null, nullif, numeric_truncation, of, off, offsets, on, once, only, open, option, or, order, over, override, perm, permanent, pipe, plan, precision, prepare, primary, print, privileges, proc, procedure, processexit, public, raiserror, read, readtext, reconfigure, references, repeatable, replace, replication, reserved_pgs, return, revoke, role, rollback, rowcnt, rowcount, rows, rule, save, schema, scroll, select, serializable, session_user, set, setuser, shared, shutdown, some, statistics, stripe, sum, syb_identity, syb_restree, system_user, table, tape, temp, temporary, text, textsize, then, time, title, to, tran, transaction, trigger, trunc, truncate, tsequal, type, uncommitted, union, unique, update, updatetext, use, used_pgs, user, user_option, using, values, varying, view, waitfor, when, where, while, with, work, write, writetext.

Note Do not use the Euro character or other non-alphanumeric or special characters, such as ' ~ ! @ # \$ % ^ & * () + - = [] \ { } | ; ` : " , . / < > ? in CS database (library system) names.

Content Services in a Cluster Environment

As of release 5.5, Content Services supports Microsoft Windows Cluster Server (MSCS) technologies under the following conditions:

- Only two-node failover clusters are supported.
- The participating cluster nodes must be running Windows 2000 Advanced Server.
- Microsoft Cluster Server must be installed on a local drive of each node in the cluster.
- Oracle Fail Safe 3.3.3 must be installed on a local drive of each node in the cluster. Oracle Fail Safe Manager may be installed on one or both cluster nodes, or on a client machine outside the cluster.
- The database engine must be Oracle9i 9.2.0.
- The database-engine node may be either in the Content Services cluster or outside it, possibly in its own cluster.
- Content Services runs in an active/passive failover configuration on a drive shared by each node in the cluster.
- All Content Services files and directories must be installed on the shared drive.
- Each Content Services library system has its own unique virtual server name, identical to the library name.
- You cannot rename a server once it becomes a member of a cluster.

Refer to the online Windows 2000 Clustering documentation or go to the [Microsoft](#) web site for the most detailed and up-to-date information on hardware requirements, system performance, and procedures for configuring clusters.

To reach the Windows Clustering documentation, choose Start > Programs > Administrative Tools > Configure Your Server. In the Windows 2000 Configure Your Server window, click Advanced and then click Cluster Service. To install and set up the cluster environment, refer to the following topics:

- [“Installing and Configuring Microsoft Cluster Server” on page 84](#)
- [“Installing Oracle Fail Safe 3.3.3” on page 85](#)
- [“Verifying Oracle Fail Safe” on page 86](#)

To bring the Oracle database into a cluster environment, see:

- [“Creating and Verifying a Database Group” on page 87](#)
- [“Creating and Verifying the Database” on page 87](#)

To bring the library systems (created during Content Services installation) online see:

- [“Controlling Library System Services in a Cluster Environment” on page 202.](#)

To uninstall Content Services in a cluster environment, see:

- [“Uninstalling Content Services in a Cluster Environment” on page 219](#)

Installing and Configuring Microsoft Cluster Server

Before beginning cluster installation, read the **Planning for Windows Clustering** “Requirements for server clusters” topic in the online help of Windows 2000 Advanced Server to be sure that you have the proper hardware required for clustering. And be sure that each node that will be in the cluster satisfies the hardware and software requirements shown in [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#).

To install Microsoft Cluster Server:

1. Install Windows 2000 Advanced Server on each node of the cluster, accepting default values for settings.
2. From a command prompt, type `comclust.exe`.
3. Repeat step 2 at each node of the cluster.

To set up the cluster:

1. Verify that all the nodes in your cluster are in the same Windows 2000 domain. To add a node to the domain:
 - a. Open the Windows 2000 Control Panel and start the System tool.
 - b. Go to the Network Identification tab and click Properties.
 - c. Select the Domain option button and enter your domain name. For example, to join the MYDOM.LOCAL domain, type MYDOM and click OK.
 - d. Reboot the server.
2. Add a user who is a member of the Domain Admins group to the local Administrators group of each node in your cluster.
3. Using Cluster Administrator, add each node to the cluster you are configuring. Refer to Windows Clustering documentation for additional information.

To prepare the Content Services environment:

1. On the Domain Controller, use DNS to assign an unused static IP address for each new library system you are going to install:
 - a. Choose Start > Programs > Administration Tools > DNS.
 - b. Right-click the name of the domain in the left pane of the DNS window and choose New Host.
 - c. In the New Host dialog box, type the name of the new library system and a unique IP address. Use this name as the database name during Content Services library system installation when prompted for the database name.

This name will also be referred to later as the *library virtual server name*, which should not be confused with the *virtual server name* or the *Oracle virtual server name*.
 - d. Click Add Host.
 - e. Repeat steps b through d for each new library system you are going to install.
2. On the Domain Controller, use DNS to assign an unused static IP address for the Verity K2 Master Admin service that will govern the library systems:

- a. Choose Start > Programs > Administration Tools > DNS.
 - b. Right-click the name of the domain in the left pane of the DNS window and choose New Host.
 - c. In the New Host dialog box, type the virtual Verity service name in the form CS_cluster_name_VERITY and a unique IP address.
 - d. Click Add Host.
3. Choose the node you want as the *cluster owner*, where you will install the Oracle database engine and the Content Services library system.
 4. Power down the non-cluster owner nodes, restart the cluster owner node, and log on to it as a user in the Domain Admins group.
 5. Restart the non-cluster owner nodes in the cluster, logging on to each as a user in the Domain Admins group.
 6. Start Cluster Administrator on the cluster owner and verify that the Cluster Owner is the node you restarted in step 4.

Continue at “[Installing Oracle Fail Safe 3.3.3](#)” on page 85.

Installing Oracle Fail Safe 3.3.3

If you have not already done so:

- Uninstall any previous version of Oracle Fail Safe from each cluster node.
- Install the Oracle software on a local drive of each node in the cluster (see “[Setting Up the Oracle Software](#)” on page 89), but do *not* create your Oracle database yet.

Note Install Oracle software on each cluster node in identical ORACLE_HOME directories, and select identical options for both installations.

- Log on to each cluster node as the same user member of the Domain Admins group for the domain to which the nodes belong.
- Stop Microsoft Cluster Administrator on each node before you start installing Oracle Fail Safe on that node.

To install Oracle Fail Safe, follow the instructions included in the online documentation found in the *Oracle Fail Safe* CD-ROM. The following constraints apply:

- Finish installing Fail Safe Server on one cluster node *before* you start installing it on the other cluster node.
- Optionally, install Fail Safe Manager on a cluster node when you are installing Fail Safe Server on that node.

You must install Fail Safe Manager on at least one cluster node or on a client node outside the cluster to verify the success of the Fail Safe installation (see “[Verifying Oracle Fail Safe](#)” on page 86)

- Install Fail Safe Server (and Manager) only on a local drive, *not* on a shared cluster drive.
- Choose Typical as the Installation Type.
- Specify the same Oracle Home for Fail Safe Server (and Manager) on each node in the cluster, and make sure that no other Oracle products are installed in the Oracle Home.

- When prompted for the Oracle Fail Safe domain, specify the domain to which the cluster nodes belong. When prompted for the account and password, specify the same user (a member of the Domain Admins group) who logged in to each node of the cluster.
- Reboot the node when you have finished installing Fail Safe Server (and Manager) if this is the first time you have installed them in the Oracle Home location. Be sure to wait until the reboot process is complete (when you see the login prompt) before you start installing Fail Safe on the other cluster node.

When you have completed the installation of Oracle Fail Safe on each node of the cluster, continue at “[Verifying Oracle Fail Safe](#)” on page 86.

Verifying Oracle Fail Safe

After installing Oracle Fail Safe, you should verify that the installation is correct by performing the following procedure. Check the installation instructions found in the documentation on the *Oracle Fail Safe* CD-ROM for details on each step.

To verify the Oracle Fail Safe installation:

1. Invoke Oracle Fail Safe Manager from the Start Menu.
2. When Fail Safe Manager displays the Connect to Cluster dialog box, type the user name, password, cluster alias, and domain to which the cluster belongs:
 - The user name (a member of the Domain Admins group) and password should be what you supplied when logging in to each node of the cluster in “[Installing Oracle Fail Safe 3.3.3](#)” on page 85.
 - The cluster alias is the name you assigned to the cluster when you created it.
3. When Fail Safe Manager displays its main window, verify that the cluster nodes are shown.
4. When connecting for the first time to a cluster after installing Oracle Fail Safe, you will be prompted to run the Verify Cluster procedure. (You can run Verify Cluster at any time by choosing Troubleshooting > Verify Cluster from the Fail Safe Manager menu bar.)

Verify Cluster checks to see that:

- Oracle Home is the same on each cluster node
 - the same version of Oracle Fail Safe is on each cluster node
 - the host name-to-IP address associations are consistent for each node in the cluster
5. On each cluster node, check to see that the Services Control Panel contains an entry called OracleFailSafe. The Startup status should be *Started* on the node where the Cluster Group resides and *Manual* on the other cluster node. (The Cluster Group is the virtual server created by Microsoft Cluster Server.)
 6. Invoke Microsoft Cluster Administrator on each cluster node and verify that Oracle Fail Safe Server is a resource in the Cluster Group.
 7. Using Microsoft Cluster Administrator, verify that the Oracle database server is displayed as an Oracle Database resource type.

If you found any errors while performing the procedure above, you will need to correct them before you can complete the installation of the Oracle server in the cluster. Otherwise, continue at [“Creating and Verifying a Database Group” on page 87](#)

Creating and Verifying a Database Group

After verifying your installation of Oracle Fail Safe, you must now create and verify a resource group in which the Oracle database server will reside.

To create and verify a resource group:

1. Invoke the Oracle Fail Safe Manager and choose Groups > Create to launch the Create Group Wizard.
2. Give the group a name and (optional) description, specify a failback policy, and a preferred cluster node for the group.
3. When prompted to add a virtual address to the group, click Yes to invoke the Add Virtual Address Wizard. (The virtual address is needed before you can add a resource, such as the Oracle database server, to the group.)
4. Using the Add Virtual Address Wizard, select the Virtual Address resource type and supply the Oracle Virtual Server name as the Host Name. The wizard should then resolve the name to an IP address.
5. In the tree view of the Fail Safe Manager, right-click the group name and choose Verify Group. Fail Safe Manager will display the progress of the operation and display a status message when the operation is complete.

Creating and Verifying the Database

When an Oracle database resides in a cluster environment, it must be installed on the drives shared by the cluster, and it must be a resource in a group. If you have not already done so, create the group for the database (see [“Creating and Verifying a Database Group” on page 87](#)).

To create and verify the database:

Create the database by following the instructions in [“Creating an Oracle9i Database” on page 102](#). Be sure to specify a shared cluster drive for the location of the database files. (This includes the control files, data files, log files and the locations determined by the LOG_ARCHIVE_DEST parameters.) Take note of the Oracle virtual server name, as you will need to supply this name during Content Services installation.

The **INIT.ORA** file can be either on the shared drives or on the local drives of the cluster nodes. If it is on the local drives, it must have the same local path on each drive.

Do *not* create any property repositories yet; you will do that after the database has been added to its resource group.

6. Oracle Fail Safe Manager will display an icon in the Standalone Resources folder of the tree view for the database you just created. To verify this database, select the icon and choose Troubleshooting > Verify Standalone Database.
7. In the Verify Standalone Database, you should see a value in each field. Type and confirm the password, and then click OK. A status message will appear when the verification is complete.
8. Add the verified standalone database to its resource group by dragging the database icon in the tree view from the Standalone Resources folder to the resource-group folder.

At this point, you should have a fully operational database that is aware of both Oracle Fail Safe and its cluster environment. To finish preparing the database for use by Content Services components, see:

- “Creating a Property Repository Tablespace” on page 106
- “Optimizing Oracle” on page 108

Setting Up the Oracle Software

Before installing Content Services 5.5.0, you need to install (or upgrade to) Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0 software, as described in the following sections.

You can install the Oracle software on a Windows platform or any of three UNIX platforms: HP-UX, or Sun Solaris. The procedures in this section provide minimal instructions for setting up your Oracle database. Details are given only where Content Services has specific requirements pertaining to installation and configuration. Refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#) for database hardware and software requirements specific to Content Services. For complete database installation and configuration requirements, refer to the Oracle documentation.

After you have set up the Oracle database, installed Content Services, and added large amounts of data to your library systems, your database administrator may occasionally need to adjust some database parameters or add database indexes to optimize performance.

If you are going to install Oracle software in a cluster governed by Microsoft Cluster Server, install the database engine on a local drive of each node in the cluster, but do not create a database yet.

To install the Oracle software on a Windows platform refer to one of the following topics:

- [“Installing Oracle9i on Windows” on page 92](#)
- [“Installing Oracle 10g on Windows” on page 97](#)

To install the Oracle software on a UNIX platform refer to the following topics:

- [“UNIX Settings for Oracle” on page 91](#)
- [“Installing Oracle9i on Windows” on page 92](#)
- [“Installing Oracle9i on UNIX” on page 94](#)
- [“Installing Oracle 10g on Windows” on page 97](#)
- [“Installing Oracle 10g on UNIX” on page 99](#)

After installing the Oracle software, refer to one of the following topics, which apply to both Windows and UNIX platforms, to create the Oracle database for Content Services:

- [“Creating an Oracle9i Database” on page 102](#)
- [“Creating an Oracle 10g Database” on page 104](#)

After creating the Oracle database, refer to the following topic, which applies to Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0 on both Windows and UNIX platforms, to create a property repository (tablespace) for each intended CS library system:

- [“Creating a Property Repository Tablespace” on page 106](#)

After you have finished the preceding tasks, you need to set up the connectivity software so that Content Services and the CS client application programs can communicate with the Oracle database. Refer to

- [“Connecting to the Oracle Database” on page 109](#)

If you are upgrading to Oracle10g 10.2.0.1.0 from an earlier version of Oracle, before you install Content Services 5.5.0, refer to one of the following topics:

- “Upgrading to Oracle 10g 10.2.0.1.0 on Windows” on page 144
- “Upgrading to Oracle 10g 10.2.0.1.0 on UNIX” on page 146

UNIX Settings for Oracle

Before installing the Oracle database on a UNIX server (HP-UX or Sun Solaris), you need to set up the required UNIX kernel parameters required for Oracle. There are no special Content Services requirements for these parameters beyond what's required for Oracle.

You also need to set the Oracle environment variables. Among these variables is `NLS_LANG`, by which you specify the locale (language, territory, and character set) to be used by the Oracle client application. With Content Services being the Oracle client application, you should set `NLS_LANG` as follows:

```
NLS_LANG = AMERICAN_AMERICA.WE8MSWIN1252
```

For other supported Oracle locales, see *Installing Content Services 5.5 into a non-English Environment* at [FileNet Worldwide Customer Support](#).

Throughout this documentation, references will be made to the following UNIX group and user definitions. These names are used as conventions only and may vary in your particular configuration:

- The *oracle* user refers to the account that owns (or will own) the installed database instance. You must use the *oracle* account when running the Oracle Universal Installer.
- The *dba* group refers to the Oracle database administrators group. The root user and the *oracle* user must be members of this group.
- The *oinstall* group refers to the group that will have permission to install/update the Oracle software. This group should be the primary group of the *oracle* user.

For information on setting UNIX kernel parameters, Oracle environment variables, and UNIX groups and users, refer to the Oracle documentation.

Installing Oracle9i on Windows

The following procedure provides minimal instructions for installing Windows Oracle9i Release 2 (9.2.0) on a Windows platform. Details are given only where Content Services has specific requirements pertaining to the Oracle installation and configuration. For complete installation information, refer to the Oracle documentation.

If you are installing the Oracle software within a cluster, perform the procedure the installation on each node in the cluster, but do not create an Oracle database yet.

Note It is strongly recommended that if you install Content Services co-resident with the database engine, that only one database (instance) be installed on that machine to reduce client connection complexity.

To install Oracle9i on a property server:

1. Insert the first *Oracle9i* CD-ROM into your CD-ROM drive and access Oracle9i Server - Autorun at **Disk1\autorun\autorun.exe**.
2. On the Available Products screen, click Oracle9i Database 9.2.0.1.0 and click Next.
3. On the Installation Types screen, click Custom and click Next.
4. The Available Product Components screen lists all components available for installation. Select the components you wish to install (see below) and click Next.

Except where noted, the following items are the bare minimum requirements for Content Services:

- Oracle9i 9.2.0.1.0
- Oracle Net Services 9.2.0.1.0
 - Oracle Net Listener 9.2.0.1.0
- Oracle9i Development Kit 9.2.0.1.0
 - Oracle Call Interface (OCI) 9.2.0.1.0
 - Oracle ODBC Driver 9.2.0.1.0 (optional)

Note Although Oracle ODBC Driver is neither certified nor used by Content Services, and need not be installed, you may want to install it anyway since it contains ODBC Test, a useful troubleshooting tool for testing ODBC connectivity.

You may also want to install other components for added flexibility or administration, such as Oracle Partitioning or Enterprise Manager Products.

5. On the Create Database screen, click No (to not create a database) and click Next. FileNet recommends that you create and configure a database after the Oracle component installation.
6. At the end of the installation, the Oracle Net Configuration Assistant will automatically launch. Clear the “Perform typical configuration” check box and click Next.
7. On the Listener Configuration, Select Protocols screen, ensure that TCP is in the Selected Protocols list and click Next.

8. On the Listener Configuration, TCP/IP Protocol screen, use the standard port number (1521), or choose your own, and click Next. Take note of this port number, as all the CS Servers (including web servers) and CS clients will need to use it.
9. On the Listener Configuration, More Listeners? screen, click No (to omit configuring another listener) and click Next. Only one listener is required for the database Content Services will use.
10. On the Naming Methods Configuration screen, choose No (to omit changing naming methods) and click Next.

You may create and configure a database instance later, either manually or via the Oracle Database Configuration Assistant (DBCA). To use DBCA to create and configure the instance, continue at [“Creating an Oracle9i Database”](#) on page 102.

Installing Oracle9i on UNIX

The following procedure provides minimal instructions for installing Oracle9i Release 2(9.2.0) on a UNIX platform: Sun Solaris, HP-UX, or AIX. The information included here is intended to serve only as an example of how an Oracle installation might be done in preparation for a library system installation. For your actual installation, you *must* refer to the Oracle documentation.

The Oracle installer requires an X Window client, and you must have your DISPLAY environment variable set to this client.

Note If you install Content Services co-resident with the database engine, it is strongly recommended that only one database (instance) be installed on that machine to reduce client connection complexity.

To install Oracle9i on a property server:

1. Log on to the property server as user *oracle*.
2. Mount the first CD-ROM of the Oracle9i installation software. (You must have root privileges to mount or unmount a CD-ROM.)
3. Start the Oracle Universal Installer with the following command:

`mount-point/disk1/runInstaller`
4. On the File Locations screen, accept the Source and Destination default values, or edit as necessary for your configuration. Make note of the Oracle Home name, as you will need it in future procedures. Then click Next.
5. On the Available Products screen, accept the default choice, Oracle9i Database 9.2.0.1.0, and click Next.
6. On the Installation Types screen, click Custom, and then click Next.
7. The Available Products screen lists all the Oracle components available for installation. Select the check boxes of the following components for installation. Except where noted, these are the bare minimum requirements for Content Services:
 - Oracle9i 9.2.0.1.0
 - Oracle Net Services 9.2.0.1.0
 - Oracle Net Listener 9.2.0.1.0
 - Oracle9i Development Kit 9.2.0.1.0
 - Oracle Call Interface (OCI) 9.2.0.1.0
 - Oracle ODBC Driver 9.2.0.1.0 (optional)

Although Oracle ODBC Driver is not used by Content Services and need not be installed, you may want to install it anyway since it contains ODBC Test, a useful troubleshooting tool for testing ODBC connectivity.

You may want to install other components for added flexibility or administration, such as Advanced Replication, Oracle Partitioning, or Enterprise Manager Products.

8. On the Privileged Operating System Groups screen, the names of the Database Administrator (SDBA) Group and the Database Operator (OSOPER) Group are listed. Typically, **dba** is the name for both of these groups. Click Next.

9. On the Create Database screen, click No to skip creation of a new database and then click Next.
10. When the Setup Privileges screen opens, leave it open and run the configuration script **\$ORACLE_HOME/root.sh** as the root user (in another Exceed window if you are doing the installation from a Windows machine).
11. At the end of the installation, the Oracle Net Configuration Assistant will automatically launch. Clear the “Perform typical configuration” check box and click Next.
12. On the Directory Usage Configuration screen, click No to defer configuration of a directory service and click Next.
13. On the Listener Configuration, Select Protocols screen, ensure that TCP is in the Selected Protocols list and click Next.
14. On the Listener Configuration, TCP/IP Protocol screen, use the standard port number (1521), or choose your own, and click Next. Make note of this port number, as all the CS Servers (including web servers) and CS clients will need to use it.
15. On the Listener Configuration, More Listeners? screen, click No to omit configuring another listener and click Next. Only one listener is required for the database Content Services will use.
16. On the Naming Methods Configuration screen, choose Yes and add HOSTNAME and click Next.
17. Local Net Service Name Configuration
 - Select “Add”, Click Next
 - Select Oracle8 release 8.0 database or service
 - Enter SID
 - Select Protocol = TCP
 - Host Name = Use the Oracle database Virtual Server Name
 - Standard port = 1521, Click Next
 - NO...do not perform test
 - Net Service Name = Virtual Server Name of the Oracle DB created on the Domain Controller
 - Click Next, Select “No” to configuring another Net Service Name
 - Click Finish

Note Open sqlnet.ora file and comment out (add the # signs) the first two lines.

```
#NAMES.DEFAULT_DOMAIN
#SQLNET.AUTHENTICATION_SERVICES
```

18. On the End of Installation screen, click Installed Products to verify that Oracle9i 9.2.0 is installed, then click Close.
19. In the **.profile** file of the **oracle** user, add/edit the following two entries, as needed:

```
export LD_LIBRARY_PATH=$ORACLE_HOME/lib32  
export LD_LIBRARY_PATH_64=$ORACLE_HOME/lib
```

20. Activate the revision to the **.profile** file by running the following command from the home directory of the **oracle** user:

```
. ./profile
```

Continue with “Creating an Oracle9i Database” on page 102.

Installing Oracle 10g on Windows

The following procedure provides minimal instructions for installing Windows Oracle 10g Release 2 (10.2.0.1.0) on a Windows platform. Details are given only where Content Services has specific requirements pertaining to the Oracle installation and configuration. For complete installation information, refer to the Oracle documentation.

If you are installing the Oracle software within a cluster, perform the procedure the installation on each node in the cluster, but do not create an Oracle database yet.

Note It is strongly recommended that if you install Content Services co-resident with the database engine, that only one database (instance) be installed on that machine to reduce client connection complexity.

To install Oracle 10g on a property server:

1. Insert the first Oracle 10g CD-ROM into your CD-ROM drive and access Oracle 10g Server - Autorun at autorun\autorun.exe. Click Install/Deinstall Products in autorun main menu.
2. On the Welcome screen, select “Advanced Installation” radio button and then click Next.
3. On the Select Installation Type screen, select Custom and click Next.
4. On the Specify Home Details screen, enter the Name and Destination Path of the new Oracle home. Click Next.
5. The Available Product Components screen lists all components available for installation. Select the components you wish to install (see below) and click Next.

Except where noted, the following items are the bare minimum requirements for Content Services:

- Oracle Database 10g 10.2.0.1.0
- Oracle Net Services 10.2.0.1.0
 - Oracle Net Listener 10.2.0.1.0
- Oracle Call Interface (OCI) 10.2.0.1.0
- Oracle Windows Interfaces 10.2.0.1.0
 - Oracle ODBC Driver 10.2.0.1.0 (optional)

Note Although Oracle ODBC Driver is neither certified nor used by Content Services, and need not be installed, you may want to install it anyway since it contains ODBC Test, a useful troubleshooting tool for testing ODBC connectivity.

You may also want to install other components for added flexibility or administration, such as Oracle Partitioning or Enterprise Manager Products.

6. On the Create Database screen, select “Install database software only” and click Next. FileNet recommends that you create and configure a database after the Oracle component installation.
7. On the Oracle Services for Microsoft Transaction Server enter the port number and click Next.
8. The Summary screen would show the summary of the component to be installed. Click Install.

9. Configure Listener and Local Net Service Name.

- Launch the Oracle Net Configuration Assistant tool.
- On Welcome screen, select “Listener Configuration” and click Next.
- Select “Add”, click Next.
- Specify “Listener Name”, click Next.
- The Selected protocols list should contain TCP as one of the values. Click Next.
- Specify the TCP/IP Port Number. Click Next.
- Select No for configuring another listener.
- Welcome screen is displayed. Select “Local Net Service Name Configuration”. Click Next.
- Select “Add”, click Next.
- Specify the Service Name (SID), click Next.
- Select TCP as protocol, click Next.
- Specify “Host Name” (use database Virtual Server name) and port number (standard port number in most cases), click Next.
- Select “No. do not test”, click Next.
- Specify “Net Service Name” (Virtual Server Name of the Oracle DB created on the Domain Controller), click Next.
- Select No for configuring another Net Service Name. Click Next.
- Click Finish.

You may create and configure a database instance later, either manually or via the Oracle Database Configuration Assistant (DBCA). To use DBCA to create and configure the instance, continue at [“Creating an Oracle 10g Database” on page 104](#).

Installing Oracle 10g on UNIX

The following procedure provides minimal instructions for installing Oracle 10g Release 2 (10.2.0.1.0) on a UNIX platform: Sun Solaris, HP-UX, or AIX. The information included here is intended to serve only as an example of how an Oracle installation might be done in preparation for a library system installation. For your actual installation, you must refer to the Oracle documentation.

The Oracle installer requires an X Window client, and you must have your DISPLAY environment variable set to this client.

Note If you install Content Services co-resident with the database engine, it is strongly recommended that only one database (instance) be installed on that machine to reduce client connection complexity.

To install Oracle 10g on a property server:

1. Log on to the property server as user oracle.
2. Mount the first CD-ROM of the Oracle 10g installation software. (You must have root privileges to mount or unmount a CD-ROM.)
3. Start the Oracle Universal Installer with the following command:

```
mount-point/runInstaller
```

4. On the Welcome screen, select “Advanced Installation” radio button and then click Next.
5. On Specify Inventory directory and credentials screen, specify the full path of the inventory directory and the operating system group name. Click Next.
6. On the Select Installation Type screen, select Custom and click Next.
7. On the Specify Home Details screen, enter the Name and Destination Path of the new Oracle home. Click Next.
8. The Available Product Components screen lists all components available for installation. Select the checkboxes of the following components for installation and click Next. Except where noted, the following items are the bare minimum requirements for Content Services:
 - Oracle Database 10g 10.2.0.1.0
 - Oracle Net Services 10.2.0.1.0
 - Oracle Net Listener 10.2.0.1.0
 - Oracle Call Interface (OCI) 10.2.0.1.0
 - Oracle ODBC Driver 10.2.0.1.0 (optional)

Note Although Oracle ODBC Driver is not used by Content Services and need not be installed, you may want to install it anyway since it contains ODBC Test, a useful troubleshooting tool for testing ODBC connectivity.

You may also want to install other components for added flexibility or administration, such as Oracle Partitioning or Enterprise Manager Products.

9. On Product Specific Prerequisite Checks screen, click Next.

10. On the Privileged Operating System Groups screen, the names of the Database Administrator (OSDBA) Group and the Database Operator (OSOPER) Group are listed. Typically, **dba** is the name for both of these groups. Click Next.
11. On the Create Database screen, select “Install database software only” and click Next. FileNet recommends that you create and configure a database after the Oracle component installation.
12. The Summary screen would show the summary of the component to be installed. Click Install.
13. After installation is complete, “Execute Configuration scripts” screen is displayed. Open a new terminal window to login as root user and run the scripts as listed. Click OK to continue.
14. Click Exit to end the installation.
15. Configure Listener and Local Net Service Name.
 - Launch the Oracle Net Configuration Assistant tool. (using netca command).
 - On Welcome screen, select “Listener Configuration” and click Next.
 - Select “Add”, click Next.
 - Specify “Listener Name”, click Next.
 - The Selected protocols list should contain TCP as one of the values. Click Next.
 - Specify the TCP/IP Port Number. Click Next.
 - Select No for configuring another listener.
 - Welcome screen is displayed. Select “Local Net Service Name Configuration”. Click Next.
 - Select “Add”, click Next.
 - Specify the Service Name (SID), click Next.
 - Select TCP as protocol, click Next.
 - Specify “Host Name” (use database Virtual Server name) and port number (standard port number in most cases), click Next.
 - Select “No. do not test”, click Next.
 - Specify “Net Service Name” (Virtual Server Name of the Oracle DB created on the Domain Controller), click Next.
 - Select No for configuring another Net Service Name. Click Next.
 - Click Finish.

Note Open the sqlnet.ora file and comment out (add the # signs) the first two lines.

```
#NAMES.DEFAULT_DOMAIN
#SQLNET.AUTHENTICATION_SERVICES
```

16. In the **.profile** file of the **oracle** user, add/edit the following two entries, as needed:

```
export LD_LIBRARY_PATH=$ORACLE_HOME/lib32
export LD_LIBRARY_PATH_64=$ORACLE_HOME/lib
```

17. Activate the revision to the **.profile** file by running the following command from the home directory of the **oracle** user:

```
. ~/.profile
```

Continue with “Creating an Oracle 10g Database” on page 104.

Creating an Oracle9i Database

The following instructions utilize the Oracle Database Configuration Assistant (DBCA) for creating an Oracle database instance. For creating an instance manually, please refer to Oracle9i documentation. The DBCA will allow you to choose specific parameter values that can affect Oracle performance. It is strongly recommended that you consult your Oracle documentation before creating a database to help you determine the best configuration for how you will use Content Services.

Note It is strongly recommended that you use a dedicated instance for Content Services. Other applications using the same instance Content Services is using can have potential conflicts and/or performance problems with Content Services, the other application(s), or both.

When using DBCA to create a database, you will need to choose a database template based on how you will be using Content Services. DBCA uses your template choice to install a pre-configured database suitable to a particular type of use:

To facilitate efficient Oracle9i 9.2.0 cost-based optimizer (CBO) processing and performance of your library system searches, you should choose the Transaction Processing template when creating the database that will contain your library system (tablespace). The Transaction Processing template is suitable for many common transactions that return moderate amounts of data, such as occurs in Content Services.

Before you create an Oracle9i database, you should have already installed an Oracle9i database engine (see [“Creating an Oracle9i Database” on page 102.](#))

To create and configure an Oracle9i database:

1. Launch the DBCA.
2. On the Operations screen click “Create a database” and click Next.
3. On the Database Templates screen, click the Transaction Processing option. You can click Show Details to see the template parameter values for this template. Then click Next.
4. On the Database Identification screen, enter a Global Database Name and an Oracle instance identifier (SID). The database name must have at most eight characters. Click Help to see additional restrictions on the database name and SID. Then click Next.

Caution Do not enter *cluster* as the Global Database Name, whether or not you have installed the database engine in a cluster environment. If you do, your database installation will fail.

5. On the Database Connection Options screen, click Dedicated Server Mode and click Next. FileNet recommends Dedicated Server Mode for better performance.
6. On the Initialization Parameters screen, you can edit the parameter values for your database. For a complete list of the parameters whose values you can edit, click All Initialization Parameters. Content Services needs the following parameter values:
 - On the Memory tab, click Custom, click All Initialization Parameters, and verify that the following parameter values in the All Initialization Parameters dialog box are as shown:
 - Set `shared_pool_size` to at least 50,000,000 bytes. The larger the value, the better the performance, subject to the size of physical memory. (The installation default value is 50,331,648 bytes.)
 - Set `db_cache_size` to 33,554,432 bytes or higher. (The installation default value is 25,165,824 bytes.)

- Set processes to 150, or the maximum number of connections + 10%, whichever is greater.
- OPTIMIZER_MODE = FIRST_ROWS
- HASH_JOIN_ENABLED = FALSE

The OPTIMIZER_MODE value causes the CBO to optimize your queries for the best response time to return the initial rows.

The HASH_JOIN_ENABLED value of FALSE turns off hash joins. Through testing Content Services for quickest response time for various queries and row set size using a large database, it has been found that the FALSE value greatly improves performance for many common operations.

Note If the oracle software version is Oracle9i 9.2.0.6, then click on Show Advanced Parameters and set BLANK_TRIMMING=TRUE.

- On the Character Sets tab, choose a character set. If you are installing Content Services in an English or other Western European language environment (such as French or German), FileNet recommends WE8MSWIN1252 as your character set since all the FileNet CS clients are Windows applications. For other installation environments, refer to the technical notice *Installing Content Services 5.5 into a Non-English Environment*, available at [FileNet Worldwide Customer Support](#).make sure that the character set is WE8MSWIN1252.
- On the DB Sizing tab, set the Sort Area Size to at least 524288 bytes.

7. On the Creation Options screen, click Create Database and click Finish.

After creating the database, you can create a tablespace to be used by a Content Services library system. Refer to “Creating a Property Repository Tablespace” on page 106.

Creating an Oracle 10g Database

The following instructions utilize the Oracle Database Configuration Assistant (DBCA) for creating an Oracle database instance. For creating an instance manually, please refer to Oracle 10g documentation. The DBCA will allow you to choose specific parameter values that can affect Oracle performance. It is strongly recommended that you consult your Oracle documentation before creating a database to help you determine the best configuration for how you will use Content Services.

Note It is strongly recommended that you use a dedicated instance for Content Services. Other applications using the same instance Content Services is using can have potential conflicts and/or performance problems with Content Services, the other application(s), or both.

When using DBCA to create a database, you will need to choose a database template based on how you will be using Content Services. DBCA uses your template choice to install a pre-configured database suitable to a particular type of use.

To facilitate efficient Oracle 10g 10.2.0.1.0 cost-based optimizer (CBO) processing and performance of your library system searches, you should choose the Transaction Processing template when creating the database that will contain your library system (tablespace). The Transaction Processing template is suitable for many common transactions that return moderate amounts of data, such as occurs in Content Services.

Before you create an Oracle 10g database, you should have already installed an Oracle 10g database engine (“[Installing Oracle 10g on Windows](#)” on page 97 or “[Installing Oracle 10g on UNIX](#)” on page 99.)

To create and configure an Oracle 10g database:

1. Launch the DBCA.
2. Click Next on the Welcome screen.
3. On the Operations screen select “Create a database” option and click Next.
4. On the Database Templates screen, click the Transaction Processing option. You can click Show Details to see the template parameter values for this template. Then click Next.
5. On the Database Identification screen, enter a Global Database Name and an Oracle instance identifier (SID). The database name must have at most eight characters. Click Help to see additional restrictions on the database name and SID. Then click Next.

Caution Do not enter cluster as the Global Database Name, whether or not you have installed the database engine in a cluster environment. If you do, your database installation will fail.

6. In the Management Options screen, check “Configure the Database with Enterprise Manager” option and “Use Database Control for Database Management”.
7. In Database Credentials screen, enter your choice.
8. In Storage Options screen, select the storage mechanism of your choice.
9. In the Database File Locations screen, check “Use Database File Location from Template”.
10. In Recovery Configuration screen, check the “Specify Flash Recovery Area”.
11. In Database Content screen, leave the default setting.

12. On the Initialization Parameters screen, you can edit the parameter values for your database. For a complete list of the parameters whose values you can edit, click All Initialization Parameters. Content Services needs the following parameter values:
 - On the Memory tab, click Custom, click All Initialization Parameters, and verify that the following parameter values in the All Initialization Parameters dialog box are as shown:
 - Set `shared_pool_size` to at least 100MB. The larger the value, the better the performance, subject to the size of physical memory.
 - Set `db_cache_size` to 32MB or higher.
 - Set `processes` to 150, or the maximum number of connections + 10%, whichever is greater.
 - Set `sort_area_size` to at least 524288 bytes.
 - `OPTIMIZER_MODE = FIRST_ROWS`

The `OPTIMIZER_MODE` value causes the CBO to optimize your queries for the best response time to return the initial rows.

Note If the oracle software version is Oracle 10g 10.1.0.4 or Oracle 10g 10.2.0.1, then click on Show Advanced Parameters and set `BLANK_TRIMMING=TRUE`.
 - On the Character Sets tab, choose a character set. If you are installing Content Services in English or other Western European language environment (such as French or German), FileNet recommends WE8MSWIN1252 as your character set since all the FileNet CS clients are Windows applications. Make sure that the character set is WE8MSWIN1252.
 - On the Connection Mode tab, set database operation mode as Dedicated Server Mode.
13. In Database Storage screen, click Next.
14. In Creation Options screen, check Create Database option and click Finish.
15. After creating the database, you can create a tablespace to be used by a Content Services library system. Refer to [“Creating a Property Repository Tablespace”](#) on page 106.

Creating a Property Repository Tablespace

If you have not already created an Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0 database on your property server, you must do so before performing the following procedure. Refer to “Creating an Oracle9i Database” on page 102 or “Creating an Oracle 10g Database” on page 104.

Before you can install a new library system, you must create a property repository tablespace in the Oracle database installed on your property server. When you install a library system in a subsequent procedure, you will use this tablespace name as your library system name. Refer to “Library System Name Requirements” on page 82 to verify that the name you assign is in compliance. You may also create a separate tablespace for indexes.

The following procedure provides minimal instructions for manually creating tablespaces using the SQL*Plus command-line utility. If you have Oracle Enterprise Manager installed and configured to manage your property server, you may prefer to use Enterprise Manager Console (for Oracle 9i 9.2.0 or Oracle 10g 10.2.0.1.0) to create tablespaces. For complete information about creating and configuring tablespaces, refer to the Oracle9i / Oracle 10g documentation.

The Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0 tablespaces you create must be locally managed.

To manually create a locally managed Oracle9i 9.2.0 tablespace, the SQL statement CREATE TABLESPACE should contain no EXTENT MANAGEMENT LOCAL clause, and the COMPATIBLE initialization parameter must be set to 9.0.0 or higher (applicable if you upgraded your database to Oracle9i 9.2.0.) / 10.0.0 or higher (applicable if you upgraded your database to Oracle 10g 10.2.0.1.0).

The Oracle 9i 9.2.0 or Oracle 10g 10.2.0.1.0 tablespaces created by Enterprise Manager Console are locally managed by default.

To create the tablespace for your library system:

1. On the property server, open a DOS command prompt window.
2. Start SQL*Plus by entering this command:

```
sqlplus
```

3. Connect to the database as the SYS user by entering this command:

```
sqlplus> sys/password@database_name as sysdba
```

4. Create the tablespace for your library system. The data file for the tablespace must be at least 1 MB if you specify AUTOALLOCATE, or at least 15 MB if you don't specify AUTOALLOCATE. For example, to create a locally managed Oracle 9i 9.2.0 or Oracle 10g 10.2.0.1.0 **100 MB** tablespace **lib1** for a library system of the same name, enter the following SQL command.

- For Oracle running on Windows:

```
create tablespace lib1
datafile 'install_drive:\[path_to_tablespace_files]\lib1.dbf' size 100m
autoallocate;
```

- For Oracle running on UNIX:

```
create tablespace lib1
datafile '[mount_point]/lib1.dbf' size 100m autoallocate;
```

Note If you wish to use the autoextend feature, after installing Content Services, use CS Admin Tools to change the following for all Property Managers:

- Repository Warning Threshold (%): 100
- Repository Full Threshold (%): 100
- 5. Optionally, create a tablespace for the indexes that will be built. The data file for the tablespace must be at least 1 MB if you specify AUTOALLOCATE, or at least 15 MB if you don't specify AUTOALLOCATE. For example, to create a locally managed Oracle 9i 9.2.0 or Oracle 10g 10.2.0.1.0 **80 MB** index tablespace **lib1x** to go with the tablespace **lib1** you created in step 4, enter the following SQL command.

- For Oracle running on Windows:

```
create tablespace lib1x
datafile 'install_drive:\[path_to_tablespace_files]\lib1x.dbf' size 80m
autoallocate;
```

- For Oracle running on UNIX:

```
create tablespace lib1x
datafile '[mount_point]/lib1x.dbf' size 80m autoallocate;
```

Optimizing Oracle

Individual parameter tuning should be done by an Oracle DBA or by using Oracle tuning utilities at runtime when your Content Services workload is known. This topic provides specific recommendations for optimizing Oracle9i or Oracle 10g database performance. For more information, see “Configuring Databases” on page 229 and the Oracle9i or Oracle 10g documentation.

Database Tuning Parameters

For the library system to run properly, your Oracle9i or Oracle 10g database parameters must be tuned as specified below.

Parameters in the init.ora File

To set these parameters, edit the **initORACLE_SID.ora** file for your database instance. Also see the comments in that file for additional information.

| Parameter | Setting |
|-------------------|---|
| processes | 150, or the maximum number of connections + 10%, whichever is greater |
| shared_pool_size | 35000000 recommended |
| log_buffer | 163840 |
| optimizer_mode | first_rows |
| hash_join_enabled | false |

The QUEUESIZE Parameter

To set the QUEUESIZE parameter, edit the address line in the **listener.ora** file as follows:

```
(ADDRESS = (PROTOCOL = TCP) (HOST = servername) (PORT = 1521) (QUEUESIZE=100))
```

Note Setting the QUEUESIZE parameter as shown above helps prevent Oracle error ORA - 12541, “No listener” from occurring when the listener backlog is too full.

Other Performance Issues

In addition to setting the initialization parameters as specified above, follow these performance guidelines:

- Run statistics routinely.
- Place redo logs and rollback segments on separate disks.
- Ensure that log switching does not occur too often; increase the size of the redo logs if it does.
- Ensure that your server has enough physical memory to accommodate the database connections.

Connecting to the Oracle Database

On each machine, other than the database server, where you are going to install library services or CS application programs (including CS Admin Tools and IDM Desktop), you must now install the Oracle Client files.

The procedures in this topic are applicable to Oracle Client connectivity, whether the Oracle database engine runs on HP, Sun Solaris, or Windows Oracle.

You must have Microsoft Data Access Components (MDAC) 2.8 installed before performing the procedures in this topic. MDAC 2.8 is available for download on the CS 5.5 CD-ROM or at [Microsoft Downloads](#).

If you have more than one Oracle Client home, use the Oracle Home Selector tool to choose one of them as your primary Oracle home to use with Content Services *before* performing the procedures in this topic.

You must also configure your database connectivity on every machine (including the database server) where you are going to install library services, or where you are going to install a CS client application. Refer to one of the following topics:

- [“Preparing Oracle9i or Oracle 10g Client Connectivity” on page 112](#)

Finally, you must create a database user that the Content Services Setup program can use to access the database. Refer to the following topic:

- [“Creating a Database User” on page 115](#)

Installing the Oracle9i Database Client Files

If your database engine is Oracle9i 9.2.0, run the following procedure on each machine (other than the database server) where you are going to install library services or a CS client application.

To install the Oracle9i Client files:

1. Insert the *Oracle9i* CD-ROM into the CD-ROM drive of the machine where you are installing Oracle9i 9.2.0 Client.
2. On the Autorun screen, click Install/Deinstall Products.
3. In the File Locations screen, select an Oracle Home name and path, and click Next.

Note If you have previously installed the Oracle8i 8.1.7 client files, you may install the Oracle9i 9.2.0 client files in the same Destination Name and Path location if you wish, but do not try to deinstall the Oracle8i 8.1.7 files after upgrading to 9.2.0.

If you wish to retain the Oracle8i 8.1.7 client software and the capability to uninstall it later, install the Oracle9i 9.2.0 client files in different path, with a different Oracle HOME name.

4. On the Available Products screen, click Oracle9i Client 9.2.0.1.0 and then click Next.
5. On the Installation Types screen, click Custom, and then click Next.
6. On the Available Product Components screen, select the following required (except as noted) components, and then click Next.
 - Oracle Network Utilities 9.2.0.1.0
 - Oracle Database Utilities 9.2.0.1.0

- Oracle Java Utilities 9.2.0.1.0
 - SQL*Plus 9.2.0.1.0
 - Oracle JDBC/OCI Interfaces 9.2.0.1.0
 - Oracle Call Interfaces 9.2.0.1.0
7. On the Component Locations screen, accept (or change as needed) the default component locations and click Next.
 8. On the Summary screen, review your choices of components to be installed. Click Install.
 9. When the files have been installed, the Oracle Net Configuration Assistant (ONCA) starts automatically. In the Welcome screen, make sure that the “Perform typical configuration” is not checked, and click Next.
 10. In the Directory Usage Configuration screen, select the check box to defer configuration for a directory service.
 11. In Select Naming Methods screen, choose Local from the Available Naming Methods list. Click the > button to move your choice to the Selected Naming Methods list, and click Next.
 12. In the Database Version screen, click “Oracle8i or later database or server” and then click Next.
 13. In the Service Name screen, type the name of the database in the Service Name box and click Next.
 14. In the Select Protocols screen, select TCP and click Next.
 15. In the TCP/IP Protocol screen, type the name of the server where the Oracle database is installed. You can use the default port number (1521) or select another port. Then click Next.
 16. In the Test screen, click Yes to perform the test to verify that the Oracle9i Client software you are installing will connect to the Oracle9i database. Click Next. If the test fails to connect, click Back to make any needed changes.
 17. In the Net Service Name screen, enter a string of up to 15 alphanumeric characters with no embedded blanks and click Next.
 18. In the Another Net Service Name? screen, click No, and then click Next.
 19. In the Net Service Name Configuration Done screen, click Next.
 20. In the Host Name screen, click Next.
 21. In the Naming Methods Configuration Done screen, click Next.
 22. In the Done screen, click Finish.
 23. In the End of Installation screen, Click Exit.

Continue at [“Preparing Oracle9i or Oracle 10g Client Connectivity”](#) on page 112.

Installing the Oracle 10g Database Client Files

Before you install the first set of initial or additional library services on a storage server where no Oracle 10g 10.2.0.1.0 database engine is installed, you must install the Oracle 10g 10.2.0.1.0 Client software. The same requirement applies if you are going to install a CS client application, including CS Admin Tools and IDM Desktop.

Note The Oracle installer requires an X Window client, and you must have your DISPLAY environment variable set to this client.

To install the Oracle 10g Client files:

1. Insert the Oracle 10g Client CD-ROM into the CD-ROM drive of the machine where you are installing Oracle 10g 10.2.0.1.0 Client.
2. Log on to the server as user oracle.
3. Access the mount point directory for the device containing the Oracle9i CD-ROM. For example,


```
cd /cdrom
```
4. Run the Oracle Universal Installer:


```
./runInstaller
```
5. On Welcome screen, click Next.
6. On Select Installation Type screen, select Administrator and click Next.
7. On Specify Home Details screen, specify Oracle Name and Path. Click Next.
8. On the Product-Specific Prerequisite checks, click Next.
9. On Summary screen, click Install.
10. When the files have been installed, the Oracle Net Configuration Assistant (ONCA) starts automatically. In the Welcome screen, make sure that the “Perform typical configuration” is not checked, and click Next.
11. In Select Naming Methods screen, choose Local from the Available Naming Methods list. Click the > button to move your choice to the Selected Naming Methods list, and click Next.
12. In the Service Name screen, type the name of the database in the Service Name box and click Next.
13. In the Select Protocols screen, select TCP and click Next.
14. In the TCP/IP Protocol screen, type the name of the server where the Oracle database is installed. You can use the default port number (1521) or select another port. Then click Next.
15. In the Test screen, click Yes to perform the test to verify that the Oracle 10g Client software you are installing will connect to the Oracle 10g database. Click Next. If the test fails to connect, click Back to make any needed changes.
16. In the Net Service Name screen, enter a string of up to 15 alphanumeric characters with no embedded blanks and click Next.
17. In the Another Net Service Name? screen, click No, and then click Next.

18. In the Net Service Name Configuration Done screen, click Next.
19. In the Naming Methods Configuration Done screen, click Next.
20. In the Done screen, click Finish.
21. After installation is complete, “Execute Configuration scripts” screen is displayed. Open a new terminal window to login as root user and run the scripts as listed. Click OK to continue.
22. In the End of Installation screen, Click Exit.

Preparing Oracle9i or Oracle 10g Client Connectivity

The CS Storage Manager, Content Search Manager, and Replication Services are essentially database clients. In order to establish a connection to the Oracle database, you need to configure a net service name on the database server. For Windows clients, you also need to configure the ODBC connection.

For Windows clients, you also need to configure the ODBC connection using Microsoft Data Access Components (MDAC) 2.8, which must be installed before you continue. MDAC 2.8 is available for download on the CS 5.5 CD-ROM or at [Microsoft Downloads](#).

You can verify that your Windows clients are using MDAC 2.8 by checking the Version properties of the file **msdadc.dll**, usually found in the **\Program Files\Common Files\System\Ole DB** folder.

If you have more than one Oracle Client home on the storage server, use the Oracle Home Selector tool to choose one of them as your primary Oracle home to use with Content Services *before* performing the procedures in this topic.

To create a Net Service Name for UNIX servers:

1. Start the Oracle Net Manager. From \$ORACLE_HOME/bin, enter this command:

```
./netmgr
```
2. In Oracle Net Configuration, expand the Local tree.
3. Expand the Service Naming folder. If there is an entry for your database, delete it.
4. Select the Service Naming folder, then choose Edit > Create.
5. In the Welcome screen, enter the database hostname for the Net Service Name and click Next.
6. In the Protocol Settings screen, select TCP/IP (Internet Protocol) for the network protocol and click Next.
7. Enter the database server hostname for Host Name. Specify a port number (the default is 1521) for the listener on the database server and click Next.
8. Choose to enter either the 8i database Service Name or the SID for Oracle 8.0 or earlier. If you choose the former, you need to know the SERVICE_NAME, found in the **initORACLE_SID.ora** file. If you choose the latter, you will need to know the ORACLE_SID name, which is part of the name of the **init.ora** file, such as **initORCL.ora**. Select Database Default for the connection type. Click Next.
9. Perform a test (your database instance and listener must be running).

The test should be successful. You may have to choose a different database login/password than the default if you get a login failure message. If you get an error regarding unable to resolve SERVICE_NAME, then make sure you entered the correct Oracle SID or SERVICE_NAME.

Click Close, then click Finish.

10. Choose File > Save Network Configuration.

11. Choose File > Exit.

To create a Net Service Name for Windows clients and web servers:

1. Start Oracle Net Manager.
2. On the Welcome screen, click Next.
3. Select Local Net Service Name configuration, and click Next.
4. Select Add, and click Next.
5. Choose either 8i database service or 8.0/7.0 database service. If you choose the former, you need to know the SERVICE_NAME, found in the **initORACLE_SID.ora** file. If you choose the latter, you will need to know the ORACLE_SID name, which is part of the name of the **init.ora** file, such as initORCL.ora. Click Next.
6. Choose TCP protocol. Click Next.
7. For hostname, enter the database server name. Enter the port number and click Next.
8. Choose Yes, perform a test (your database instance and listener must be running). Click Next.
9. The test should be successful. You may have to choose a different database login/password than the default if you get a login failure message. If you get an error regarding unable to resolve SERVICE_NAME, then make sure you chose the correct Oracle SID or SERVICE_NAME.
10. For a Net Service Name, enter a string of up to 15 alphanumeric characters with no embedded blanks. Click Next.
11. There is no need to configure more than one Net Service Name.

Note Before creating an ODBC data source name, you need to install a Microsoft registry patch to prevent an error from occurring when you connect to an Oracle 9i / Oracle 10g database using MDAC 2.8. To install the patch:

1. Copy the appropriate patch file from the CS Admin Tools CD-ROM to a local drive on the machine where you will create an ODBC data source name:
 - **mtxoci9x_win2k.reg** if the machine is running Windows 2000 or Windows XP and has Oracle 9i client.
 - **mtxoci9x_winnt.reg** if the machine is running Windows 98 or Windows NT Workstation 4.0 and has Oracle 9i client.
 - **mtxoci10x_win2k.reg** if the machine is running Windows 2000 or Windows XP and has Oracle 10g client.
2. Double-click **mtxoci9x_win2k.reg** or **mtxoci9x_winnt.reg** or **mtxoci10x_win2k.reg** to install the patch and then reboot the machine.

To create the ODBC for Windows Clients:

1. You must have created your CS tablespace before you can configure the ODBC connection.
2. From the Control Panel, double-click the Data Sources (ODBC) icon to open the ODBC Data Source Administrator dialog.
3. Click the System DSN tab in the ODBC Data Source Administrator.
4. Click Add.
5. Double-click Microsoft ODBC for Oracle to bring up the Oracle8 ODBC Driver Setup dialog box.
6. Click Options to expand the dialog box.
7. Enter the Data Source Name (use the Net Service Name), the Server name, and check the SQLDescribeCol always returns a value for precision check box.
8. Click OK.
9. If you wish, you can test the ODBC connection using the Oracle ODBC Test utility. Be sure to select the Data Source Name for the machine you just created.

To test your Oracle connection:

1. Enter the following command in SQL*Plus to connect to the database instance:

```
sys/sys_pwd@service_name as sysdba
```

2. If you receive an error, make sure that you have entered the correct information for your Local Net Service Name configuration and for the ODBC connection. Review your **initORACLE_SID.ora**, **listener.ora**, and **tnsnames.ora** files. Also make sure that the listener service is started.

Creating a Database User

Before you can install a new library system, you must create a corresponding database user, which is used by the Content Services Setup program. You can create the user via a SQL script—your own or the script **loginscript.sql** in the *Content Services* CD-ROM. Alternatively, use the Oracle Security Manager to create the database user.

Note If you have not already created a property repository tablespace in your Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0 database, create one now before running the script.

Creating a Database User via a SQL Script

You can use the script **loginscript.sql** (see below) in the *Content Services* CD-ROM, or you can create your own. To run the script, you must be logged on as the SYS user with SYSDBA privileges. Since each library system must have its own database user, you must run the script once per library system.

Before running **loginscript.sql**, substitute your library system (tablespace) name and password in place of LIBRARYNAME and LIBRARYNAMEPWD. (The password cannot be null.) If your library system will use a separate index tablespace, uncomment the last line of the script and substitute the index tablespace name in place of INDEXTABLESPACE.

For Oracle 9i 9.2.0

```
create user LIBRARYNAME identified by LIBRARYNAMEPWD;
grant connect to LIBRARYNAME;
alter user LIBRARYNAME default tablespace LIBRARYNAME;
alter user LIBRARYNAME quota unlimited on LIBRARYNAME;
grant create procedure to LIBRARYNAME;
grant create trigger to LIBRARYNAME;
grant select on v_$session to LIBRARYNAME;
grant select on v_$sga to LIBRARYNAME;
grant select on v_$log to LIBRARYNAME;
grant select on dba_data_files to LIBRARYNAME;
grant select on dba_free_space to LIBRARYNAME;
-- alter user LIBRARYNAME quota unlimited on INDEXTABLESPACE;
```

For Oracle 10g 10.2.0.1.0

```
create user LIBRARYNAME identified by LIBRARYNAMEPWD;
grant connect to LIBRARYNAME;
alter user LIBRARYNAME default tablespace LIBRARYNAME;
alter user LIBRARYNAME quota unlimited on LIBRARYNAME;
grant create procedure to LIBRARYNAME;
grant create trigger to LIBRARYNAME;
grant alter session to LIBRARYNAME;
grant create cluster to LIBRARYNAME;
grant create database link to LIBRARYNAME;
grant create sequence to LIBRARYNAME;
grant create synonym to LIBRARYNAME;
grant create table to LIBRARYNAME;
grant create view to LIBRARYNAME;
grant select on v_$session to LIBRARYNAME;
grant select on v_$sga to LIBRARYNAME;
grant select on v_$log to LIBRARYNAME;
grant select on dba_data_files to LIBRARYNAME;
```

```
grant select on dba_free_space to LIBRARYNAME;
-- alter user LIBRARYNAME quota unlimited on INDEXTABLESPACE;
```

The database user you create by running **loginscript.sql** has the same name as the tablespace. You will be prompted for this user name (and password) during installation of the initial library system services.

Creating a Database User via Security Manager

The following procedure for Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0 show how to create a database user using the Oracle Security Manager.

To create a database user in Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0:

1. Start the Oracle Enterprise Manager Console.
2. In the tree view choose Network > Databases > *server_name* > Security, where *server_name* identifies the machine where your database engine is installed.
3. Select and right-click Users, select Create. The Create User dialog box opens.
4. In the General tab specify the following information:
 - a. In the Name box, type the CS user name, which must be identical (including case sensitivity) to the tablespace name.
 - b. Type and confirm the password, and clear the Expire Password Now check box.
 - c. In the Tablespaces group, type the tablespace name in the Default box, and TEMP in the Temporary box.
 - d. Click Unlocked.
5. In the Role tab, select Connect from the Available list and move it to the Granted list.
6. In the System tab, select the following from the Available list and move them to the Granted list.

For Oracle 9i 9.2.0

```
CREATE PROCEDURE
CREATE TRIGGER
```

For Oracle 10g 10.2.0.1.0

```
CREATE PROCEDURE
CREATE TRIGGER
ALTER SESSION
CREATE CLUSTER
CREATE DATABASE LINK
CREATE SEQUENCE
CREATE SYNONYM
CREATE TABLE
CREATE VIEW
```

7. From the tree view in the Object tab, expand Schemas > Sys > Views. For each of the following objects, move the object to the Granted list and assign the SELECT privilege to it from the Available Privileges list:
 - V_\$SESSION

- V_\$SGA
- V_\$LOG
- DBA_DATAFILES
- DBA_FREE_SPACE

8. In the Quota tab, select the tablespace you specified in step 3, click Unlimited, and then click Create.

Procedure for Moving Indexes to a Separate Tablespace

Oracle ships a script ('indextblscript.sql') that moves existing indexes to the separate tablespace reserved for them. Use the following procedure to move the existing indexes to a separate tablespace:

```
-- Create your index tablespace and grant quota as described in the Installation
Guide.
-- Substitute the name of your library for the LIBRARYNAME placeholder in the script.
-- Substitute the index tablespace name for the INDEXTABLESPACE placeholder in the
script.
-- Login as SYS user with sysdba privileges and execute this script with sqlplus.
```

A new script with the name 'alterIndexTbsp.sql' would be created.

```
-- Inspect the resulting script for any errors
-- Be sure the library user is granted quota to the index tablespace
-- Be sure the index tablespace was created large enough for the new indexes
-- Execute the generated script 'alterIndexTbsp.sql' as the SYS user
```

This script would rebuild all indexes into the separate tablespace and update the 'c_prop_indx_tbsp' column in the CATALOG table with the name of the separate index tablespace.

Configuring for WANs and Firewalls

If you use a Wide Area Network (WAN) or a firewall that restricts port-level configurations, you can override CS's dynamic port allocation by adding entries for the Storage Manager and Content Search Manager into the **services** file.

Note Port configuration should be done before you install CS library services.

Configuring the Storage Manager's Port Number

The port number for the Storage Manager must be a number from 1024 to 65535 (inclusive) that is not already assigned in the services file. This port number is entered in the following location:

| OS Platform | Location |
|-------------|--|
| Windows | File: %SystemRoot%\system32\drivers\etc\services Entry: filenet\LibrarySystem\brserver <i>port_number/tcp</i> |
| UNIX | File: /etc/services Entry: filenet/librarysystem/brserver <i>port_number/tcp</i> |

Note If you have upgraded Content Services from a “saros” directory structure, replace “filenet” with “saros” in the location entries shown in the table above:

```
saros\librarysystem\brserver    port_number/tcp
saros/librarysystem/brserver    port_number/tcp
```

If you have upgraded Content Services from an “idmds” directory structure, replace “filenet” with “idmds” in the location entries shown in the table above:

```
idmds\librarysystem\brserver    port_number/tcp
idmds/librarysystem/brserver    port_number/tcp
```

Configuring the Content Search Manager's Port Number

The TCP/IP port numbers for a Content Search Manager of a library system are associated with the Verity components and must be entered into the services file as follows:

```
CS_server_K2Admin      8105/tcp          #Verity K2 Admin port
CS_lib_K2Broker        8106/tcp          #Verity K2 Broker port for lib
CS_lib_K2Server        8107/tcp          #Verity K2 Server port for lib
```

The *server* is the local host name (in a non-cluster environment) or the cluster name (in a cluster environment), and *lib* is the name of the library system:

Caution To change these port-number assignments after installing a CS library system, you must uninstall the Verity components, change the port-number assignments in the services file, and then reinstall the Verity components.

The service file is in the following platform-dependent location:

| OS Platform | Location |
|-------------|---|
| Windows | File: %SystemRoot%\system32\drivers\etc\services |
| UNIX | File: /etc/services |

Installing Content Services

Once you have understood the “[Prerequisites to Installing Content Services](#)” on page 80, and completed any associated prerequisite tasks (such as installing the database and creating the ODBC data source name), you are ready to install the Content Services software.

Caution If you want to have IDM Desktop or some other client application program on a server where you are going to install Content Services software, you must install the client application first. If you want CS Admin Tools to run on the same server where you are going to install CS software, you must install IDM Desktop, the CS software, and CS Admin Tools in that order. If after installing CS software you need to install or reinstall a client application on the same server, you must install the CS Client Library Files after (re)installing the client application (refer to “[Client Library Files](#)” on page 207).

You must run Content Services Setup, the CS installation program, on the machine where you want the library system services to reside. This machine may be the server where you installed the database or it may be another server. The choice is yours. Refer to “[Installation Overview](#)” on page 77 for information on where you can install library services.

If you are installing a library system on a machine that is already acting as a storage server or database server for one or more library systems, do a full backup of all those existing library systems before beginning installation of the new library system. Refer to “[Backing Up the Library System](#)” on page 336 for procedures.

Installation Settings

The information in the following table applies to the CS Setup program that installs library services. Because some of the installation settings require resources or configurations that may be difficult or impossible to supply unless you prepare them in advance, FileNet recommends that you determine the values you want to use before beginning the actual library system installation.

| Installation Prompt | Supplied Value |
|---|----------------|
| <p>Display Language</p> <p>The language you want to use for displaying property names in CS Explorer, message text sent to the Messages properties, and selectable system-supplied values (for example, Yes/No). Currently, Content Services ships with support files for Chinese, English, French, German, Japanese, Korean, and Spanish (chinese.nls, english.nls, french.nls, german.nls, japanese.nls, korean.nls, and spanish.nls). Other language files can be created with direct FileNet assistance.</p> | |
| <p>ODBC Data Source Name</p> <p>The name you specified when you created an ODBC data source to refer to the database server. This name cannot be changed during library system installation.</p> | |
| <p>Tablespace Name</p> <p>The name of the Oracle tablespace you created for the library system's property repository. Refer to “Library System Name Requirements” on page 82 for information about tablespace naming requirements.</p> | |

| Installation Prompt | Supplied Value |
|--|----------------|
| <p>Temporary Library System Password</p> <p>The password for the Content Services user you added. Setup needs this password to log onto the database to configure your library system.</p> <p>Refer to “Creating a Database User” on page 115 for information about how to specify this password.</p> | |
| <p>Index Tablespace Name</p> <p>The name of the Oracle tablespace you created for the library system's index.</p> <p>Refer to “Library System Name Requirements” on page 82 for information about tablespace naming requirements.</p> | |
| <p>Master Admin Fully Qualified Server Name</p> <p>The fully qualified name (<i>servername.domain</i>) of the server where a Verity K2 Master Admin is installed (<i>servername</i> can be at most 15 characters long). A Master Admin administers the Verity components for one or more library systems.</p> <p>CS Setup will install a Master Admin on the server where you are installing the first set of Content Services components in your organization.</p> <p>If you are installing the initial services of a library system on a server where no other initial or additional services for any other library system are installed, you can have CS Setup install a Master Admin on the server or point to an existing Master Admin on another server.</p> <p>If you are installing initial or additional services of a library system on a server where services for other library systems have already been installed, CS Setup will point the newly installed CS components to the Master Admin that governs the CS components previously installed on the server.</p> | |
| <p>Master Admin Server Port Number</p> <p>The port number assigned to the Verity K2 Master Admin you have specified.</p> | |
| <p>Locale</p> <p>The language (English, French, German, etc.) and its characteristics (date formats, stemmer, etc.) that apply to all Verity collections (indexes) of a library system on the server where you are running a Content Services installation. You can specify a different locale for the Verity collections of this library system that you install on another server. Verity collections for different library systems on the same server need not have the same locale.</p> | |

| Installation Prompt | Supplied Value |
|--|----------------|
| <p>Administrative User (default: Admin)</p> <p>The name you want to assign to the initial administrative user for your library system. You will use the administrative user account to access your library system for the first time from CS Explorer. The name Admin is offered as a default. You cannot change this name after the library system is installed.</p> | |
| <p>Administrative Password</p> <p>The password you want to assign to the initial administrative user for your library system. Your password must have between 1 and 32 characters.</p> | |
| <p>Administrative Group (default: Administrators)</p> <p>The name you want to assign to the administrative group for your library system. The users you will assign to this group will be designated library system administrators. The name Administrators is offered as a default. You cannot change this name after the library system is installed.</p> | |
| <p>Universal Group (default: General Users)</p> <p>The name you want to assign to the universal group for your library system. To be a user of this library system, a person must be a member of this group. The name General Users is offered as a default. You cannot change this name after the library system is installed.</p> | |
| <p>Default Storage Category (default: Documents)</p> <p>The name you want to assign to the default storage category CS Setup will create. This storage category points to the initial storage repository that Setup creates for your library system. The name Documents is offered as a default. You cannot change this name after the library system is installed.</p> | |
| <p>Default Archive Category (default: Archive)</p> <p>The name you want to assign to the default archive category CS Setup will create. This archive category is installed in the initial reclaimable archive repository that Setup creates for your library system. The name Archive is offered as a default. You cannot change this name after the library system is installed.</p> | |
| <p>Default Document Class (default: General)</p> <p>The name you want to assign to the default document class Setup will create. The name General is offered as a default. You cannot change this name after the library system is installed.</p> | |

| Installation Prompt | Supplied Value |
|---|----------------|
| <p>CS Shared Files</p> <p>The name of directory where you want to install Content Services shared library files. The directory C:\Program Files\FileNet\Shared is offered as a default.</p> <p>Caution Do not specify a network-shared directory. If network connectivity is lost, Content Services components that depend on the shared files will fail.</p> | |
| <p>General Services</p> <p>The drive/partition where the directory containing miscellaneous Content Services components are to be installed. The free space (in MB) is shown for each available partition. You must choose a local drive/partition with at least 2.2 GB of free space. The directory will be named filenet.</p> | |
| <p>Storage Manager</p> <p>The drive/partition where the directory containing storage repository files is to be installed. The free space (in MB) is shown for each available partition. You must choose a local drive/partition with at least 20 MB of free space. For best performance, install the storage repository on its own NTFS disk/partition. The directory will be named filenet.</p> <p>Caution Do not attempt to install the storage repository on a removable disk or tape device. Serious errors will result if the installer tries to write to a device that has been removed.</p> | |
| <p>Content Search Manager</p> <p>The drive/partition where the directory containing Verity collections (indexes) for the library on this server is installed. The free space (in MB) is shown for each available partition. You must choose a local drive/partition with at least 20 MB of free space. For best performance, install the content search repository on its own NTFS disk/partition. The directory will be named filenet.</p> <p>Caution Do not attempt to install the content search repository on a removable disk or tape device. Serious errors will result if the installer tries to write to a device that has been removed.</p> | |
| <p>Installer Name</p> <p>Your name (up to 32 characters) as the installer. “Administrator” is offered as a default. In CS Explorer, this name will appear as the value of the Installer property in the System object.</p> | |
| <p>Installation Comment</p> <p>An installation comment (up to 70 characters). “Content Services 5.5” is offered as a default. In CS Explorer, the text of this comment will appear as the value for the Comment property in the System object.</p> | |
| <p>Set library system services to autostart at reboot.</p> | |

Running Content Services Setup

The Content Services Setup program (**setup.exe** on the Content Services CD-ROM) allows you to perform the following tasks:

- Installing initial library system services
- Installing additional library system services
- Reinstalling existing CS executables and related files for an existing library system
- Installing an additional display language for an existing library system

Each of these tasks is described in the following sections.

Notes

- The Content Services Setup program will not launch if you specify its location via the uniform naming convention (UNC), such as `\\host\share_name\directory\...\directory\setup.exe`.
- Assign the Windows TMP and TEMP system environment variables to partitions with sufficient disk space for the CS Setup program to ensure that you can launch the program as expected. Refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#) for disk-space requirements.

Installing Initial Library System Services

Use the following procedure to install an initial set of library system services on a server. Read through the entire procedure before beginning the installation, as some of the values you need to collect and enter require either resources or configurations that may be difficult or impossible to supply unless you prepare them in advance. If you completed “[Installation Settings](#)” on [page 119](#) as recommended earlier, then you already know what these values represent and are ready to begin.

Note If you have a previous version of a library system installed and want to upgrade to the current version, do not use this procedure. Refer to “[Upgrade Roadmap](#)” on [page 75](#) for a list of required procedures.

To install initial library system services:

1. On the database server, verify that Oracle is running:
 - a. If Oracle is installed on a Windows server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services is started.
 - b. If Oracle is installed on a UNIX server, see that the Oracle processes are running by the following command:


```
ps -eaf | grep ora
```
2. Shut down any applications that use ODBC to connect to the database, such as Enterprise Manager or Query Analyzer.
3. On the server where you are installing initial library system services (the initial storage server), stop the services for all previously installed library systems.

4. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the machine that will become the initial storage server and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.

Note (FNDTS00071809) If you are performing this installation within a Microsoft Windows Clustering environment, launch the CS Setup program from a command line as follows:

```
setup -log # !logfile.log @ALL
```

where *logfile* is the name of the installation log file you want CS Setup to generate. Due to a bug, launching CS Setup within a clustering environment by double-clicking **Drive:\Setup.exe** fails to generate a log file (**guisetup-username.log**, where *username* is the user running CS Setup).

5. On the Welcome screen, click Next.
6. On the Software License Agreement screen, click Approve, and then click Next.
7. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. If your database software is Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0, then click Next. If it isn't, click Cancel to exit Content Services Setup and refer to [“Setting Up the Oracle Software” on page 89](#).
8. In the Select Setup Operation screen, click Install Content Services, and then click Next. If you are running CS Setup in a cluster environment, continue at [step 9](#); otherwise, skip to [step 10](#).
9. On the Specify Virtual Server Name screen, enter the virtual server name (VSN), assigned by the domain administrator, that identifies the cluster where you are installing initial services. This VSN must have a DNS-assigned static address.
10. In the Specify Database Information screen, supply the requested information and click Next:
 - a. Display Language: the language in which Content Services displays property names, field names and system messages.
 - b. ODBC Data Source Name: the name you specified when you created an ODBC data source to refer to the database server. If the database is in a cluster, the ODBC data source name is the virtual database server name.
 - c. Tablespace Name: the name of the Oracle tablespace you created as the property repository for this library system.
 - d. Temporary Library System Password: The password for the Content Services user you added (refer to [“Creating a Database User” on page 115](#)). Setup needs this password to log onto the database to configure your library system.
 - e. Index Tablespace Name (optional): The name of the Oracle tablespace you created for the indexes in your library system. Select the “Use separate index tablespace” check box.
11. On the Specify Verity Settings screen, specify the Verity K2 Master Admin server name and port (as needed) and the Verity locale (language), as described below, and then click Next.
 - a. If the screen has no check box, skip to step b, since the library services you are installing will be governed by an existing Master Admin (whose server name and port are shown on the screen) either on this server or some other server. If the screen does have a check box, do one of the following:

- Select the check box and specify the fully qualified server name and port of an existing Master Admin on some other server to govern the library services you are installing.
- Clear the check box to install a Master Admin (the server name and port are shown on the screen) with the library services you are installing on the server.
- If you are installing library services in a cluster, clear the checkbox and type the fully qualified cluster name (`cs_clustername_verity.domain`) and port.

Note If no previously installed set of library services remains on this server, and no library services are installed on any other server in your organization, you *must* clear the check box so that a Master Admin will be installed with the library services you are installing on the server.

- b. The locale you choose (from the drop-down list) applies only to the Verity collections (indexes) on this server for this library system, and can differ from the locale for services of this library system on other servers. It can also differ from the locale for services of other library systems on this server.

12. In the Specify Initial Values screen, supply the requested information and click Next:

- a. Administrative User: the name of the user with access rights to every object in the library system The suggested name is Admin.
- b. Admin Password: the password of the administrative user.
- c. Administrative Group: the name of the group with access rights to every object in the library system. The suggested name is Administrators.
- d. Universal Group: the name of the group to which all library system users must belong. The suggested name is General Users.
- e. Default Storage Category: the name of the storage category that points to the initial storage repository created for this library system during installation. The suggested name is Documents.
- f. Default Archive Category: the name of the archive category that points to the initial reclaimable archive repository created for this library system during installation. The suggested name is Archive.
- g. Document Class: the name you want for the document class assigned to documents added to the library system. The suggested name is General.

13. In the Specify Install Locations screen, supply the requested information and click Next:

- a. CS Shared Files: the directory that will contain CS shared files. The suggested location is **InstallDrive:\Program Files\FileNet\Shared**. You can choose the location only if the shared files have not been previously installed on the server. Choose a location with at least 20 MB of free space. If you are installing library services within a cluster, the partition must be located on a shared drive accessible to all nodes in the cluster.

Caution Do not specify a network-shared directory. If network connectivity is lost, Content Services components that depend on the shared files will fail.

- b. For each of the following CS components, choose a partition from the drop-down list. The list will show only partitions having at least 200 MB of free space, the amount required for the CS installer. In a cluster environment, you need to choose a partition on a drive shared by all nodes in the cluster. Within the partition the CS component will be installed in the **filenet** directory.

- General Services: the partition where miscellaneous CS utilities will be installed.
- Storage Manager: the partition where the initial Storage Manager will be installed.
- Content Search Manager: the partition where the initial Content Search Manager will be installed.

Note For optimum performance, install the Content Search Manager in its own partition.

Note For General Services, choose a partition with at least 2.2 GB of free space.

After you click Next, CS Setup will prompt you to create the Shared directory if it hasn't already been created by a previous execution of CS Setup on this server. CS Setup will also prompt you to create a directory (named **filenet**), if it does not already exist, on each partition you selected for General Services, Storage Manager, and Content Search Manager. Click Yes for each prompt.

14. In the Specify Final Setup Values screen, supply the requested information and click Next.
 - a. Installer Name: the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
 - b. Installation Comment: the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is "Content Services 5.5".
 - c. Select or clear the check box that determines if you want the services for this library system to automatically start whenever the server is rebooted.

Note Occasionally Windows Oracle library systems do not automatically start after a reboot. To check the status of a library, and to manually start it, refer to "[Content Services Manager](#)" on page 198.

- d. Select or clear the check box that determines if the services for this library system will automatically start when CS Setup exits. (This check box appears only if you have already installed initial or additional services of at least one other library system on this server.)

15. The Review Setup Entries screen shows your installation selections in the preceding steps. (You can also see your installation selections in the **setup.inf** file, located at *InstallDrive:\filenet*.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and installing software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

16. In the Installing Components screen, review the results of the installation and click Next. If you are installing the first set of library services on this server, skip to [step 19](#). If you have previously installed initial or additional library system services on this server, continue at [step 17](#).

17. On the Setup Complete screen, select or clear the check box that determines whether you want to perform another setup task within the current execution of CS Setup. Then click Next. If you selected the check box, continue at [step 18](#); otherwise, skip to [step 19](#).
18. On the Select Setup Operation screen, click the option corresponding to the installation task you want to do and click Next:
 - a. To install the initial services for another library system on this server, continue at [step 9](#) (in a cluster environment) or [step 10](#) (in a non-cluster environment) of this procedure.
 - b. To install additional services on this server for a library system whose initial services are on another server, continue at [“Installing Additional Library System Services” on page 130](#).
 - c. The Reinstall Executables option is generally undertaken only as part of a recovery or move operation, such as moving the database or a set of library services to another server.
 - d. To install another display language for a library on this server, continue at [“Adding a Display Language” on page 134](#).

Note If you are installing library services in a cluster, skip the remainder of this procedure and continue at [“Configuring a Library System Within a Cluster” on page 127](#).

19. The appearance of the Exit Setup screen depends on whether you previously installed initial or additional library services on the server prior to the current execution of CS Setup:
 - a. If you previously installed library services on the server, simply click Finish to exit CS Setup.
 - b. If you are installing the first set of initial library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

After CS Setup exits (and following a restart, if you so specified) perform the steps in [“Starting and Monitoring Library Services” on page 128](#) to start the library services and to enable remote monitoring of them.

View the **cssetup.log** file, located at *InstallDrive:\filenet* by default, to see a log of the individual actions (file copying/deleting, component installation, custom property creation, etc.) in the installation of the initial library services.

View the file **GUISetup-username.log** file, located at *%TEMP%* (where *%TEMP%* is an environment variable for *username*, the user installing Content Services on this server) to see a log of the screen-related actions of CS Setup in the installation of library services.

To determine whether the Verity components installed with the initial library services are correctly configured, and to repair the configuration if necessary, run the utility CSK2CFG. Refer to [“CSK2CFG” on page 386](#) for more information.

Configuring a Library System Within a Cluster

For each initial and additional set of library services installed within a cluster, you must configure the services within the cluster, as shown in the following procedure before you start the library services. You can do this

following each installation of library services or once after installing all sets of initial and additional library services. In the procedure, *libraryname* is the name of the library system whose services you are configuring within the cluster.

To configure a library system within a cluster:

1. Shut down the non-Cluster Owner node.
2. Reboot the Cluster Owner node and log on as a user in the Domain Admins group.
3. Restart the non-Cluster Owner node and log on as a user in the Domain Admins group.
4. Start Cluster Administrator on the Cluster Owner. The Open Connection to Cluster dialog box should appear (same as File > Open Connection). Select the cluster name from the dropdown list and click Open.
5. On the left pane, expand the Groups folder icon. Then select the Cluster Group icon in the left pane. In the right pane you should verify that the Cluster IP Address, Cluster Name, and shared disk drives are owned by the node on which you installed Content Services. If they are not, use Cluster Administrator to change their owner to the Cluster Owner.
6. In the left pane, click the library system you have installed. In the right pane, right-click the IP_*libraryname* icon, and choose Properties from the pop-up menu and select the Parameters tab.
7. In the Network dropdown list, choose the entry from the Network dropdown list that corresponds to the external public LAN. Then click OK.
8. In the right pane, right-click the SHF_*libraryname*, choose Properties from the pop-up menu and select the Advanced tab.
9. Set the Threshold value to 20. (This is the number of times the Cluster services will attempt to restart the SHF_*libraryname* resource before the CS_*libraryname* group fails over. Set the “Pending timeout” value to 300 seconds. (This is the amount of time the resource in a pending state has to resolve its status before the Cluster service puts the resource into the failed state.)
10. On each node in the cluster, open the Windows Services panel and do the following for the Storage Manager of the set of library services you are configuring:
 - a. In the right pane, right-click the Storage Manager icon (CS Storage Manager - *library name*), choose Properties, and select the Recovery tab.
 - b. Assign the value “Restart the Service” to the “First failure”, “Second failure”, and “Subsequent failures” parameters.
 - c. Click OK.

Continue at “[Controlling Library System Services in a Cluster Environment](#)” on page 202 for instructions on starting the library system services within a cluster.

Starting and Monitoring Library Services

The following post-installation steps are necessary to finish configuring the library system and checking the database connection. Step 1 is necessary only when you are installing the initial services of the first library on the server.

1. To enable remote monitoring and control of library system services, the Domain Admins and Enterprise Admins groups must belong to the local Administrators group on your storage server(s). To add these groups, log on to your initial storage server as a user with at least local Administrator rights and do the following:
 - a. Right-click My Computer and select Manage from the menu.
 - b. Choose Local Users and Groups > Groups > Administrators.
 - c. On Administrators Properties, click Add.
 - d. On Select Users or Groups, select your domain from the Look In list.
 - e. Select Domain Admins and click Add, then select Enterprise Admins and click Add.
 - f. Click OK on Select Users or Groups, then click OK on Administrators Properties to add the groups.

2. Do one of the following to start your newly installed library system and register its services with Active Directory:
 - a. Log on to the initial storage server as a user with Domain Admin or Enterprise Admin rights. This is necessary so Content Services Manager will be able to create new Active Directory objects on the domain controller.
 - b. Display the Content Services Manager applet, available via the library icon in the System Tray at the far right of the Windows Taskbar, or by executing *Drive:\filenet\dsstatus.exe*.

Note If you installed a set of library services on a server, the Content Services Manager icon will not be visible until after you restart the server.
 - c. If the Service Manager does not show all of the local library system services as started, start them now.

3. Refer to “[Controlling Library System Services](#)” on page 198 for details about controlling library system services, locally or throughout an Active Directory forest.

Note The first time Content Services Manager starts a library system, it creates an Active Directory object on the domain controller for each of the library system services. In the Active Domain hierarchy, these objects appear under your server's Computer object as:

 - **SHF_LibrarySystem** for Storage Managers
 - **CSMGR_LibrarySystem** for Content Search Managers
 - **REPL_LibrarySystem** for Replication Daemons

If you experience problems during installation, refer to “[Problem Solving](#)” on page 215.

What to Do Next

Having installed the initial services of a library system, you can now proceed in several directions:

- Install (or upgrade) client software to access and administer the newly installed library system (refer to “[Accessing and Administering a Library System](#)” on page 130).
- Install (on a different server) additional services for the library system (refer to “[Installing Additional Library System Services](#)” on page 130).

- Add a display language for the library system (refer to “Adding a Display Language” on page 134).
- Run CS Setup again (on this or another server) to install the initial services of another library system (refer to “Installing Initial Library System Services” on page 123).
- Install the LDAP shared libraries. This installation is mandatory if and only if you are going to import users and groups from an LDAP directory server into a library system (refer to “LDAP_EXP” on page 435 for installation instructions).
- Install Rendition Services software (refer to “Prepare a Server” on page 662 for installation instructions).

Accessing and Administering a Library System

To access and administer the Content Services server components for a library system, you must:

- Install (or upgrade) CS Admin Tools to be able to administer the library system. A desktop and web version of this product component are available. You can administer multiple library systems from a single instance of CS Admin Tools. For installation instructions, see the **Admin\Documentation** or **WebAdmin\Documentation** directory on the *CS Admin Tools* CD-ROM.
- Install (or upgrade) the client applications accessing the library system. See “Client Interface Setup Roadmap” on page 74 for a list of associated FileNet clients.
- If you are using eProcess, verify that you can upgrade or apply the necessary hotfix pack to run with this version of Content Services. See “Client Interface Setup Roadmap” on page 74 for release number details and check the latest eProcess hotfix pack readme for notes on certification with Content Services. See the eProcess product documentation for installation instructions.

If eProcess is certified and being used at your site, verify that the following objects have been added to the library system. You can see these objects using CS Admin Tools once you have it installed.

- WorkflowDaemon user in the Admin group. The WorkflowDaemon user initially has the same password as the Admin user, which you specified when running CS Setup to install initial library services. You can use CS Admin Tools to change the password of the WorkflowDaemon user.
- IdmVWVersion internal custom property.
- Workflow Definition document class with properties Document Class and IdmVWVersion.

Installing Additional Library System Services

If you have installed the initial services for a library system on one server, you can use the following procedure to install an additional set of library system services for the library. The server where you install the additional services, the *additional storage server* for this library, cannot be the same server where you installed the initial services.

If you have a previous version of additional library services installed on the server and want to upgrade the services to the current version, do not use this procedure. Refer to “Upgrading Content Services” on page 152.

Refer to the [CS 5.5.0 Windows Oracle Platform Requirements](#) document at [FileNet Worldwide Customer Support](#) for the latest hardware and software requirements for any intended additional storage server.

Read through the entire procedure before beginning the installation of additional services, as some of the values you need to collect and enter require either resources or configurations that may be difficult or impossible to supply

unless you prepare them in advance. If you completed “[Installation Settings](#)” on page 119 before you installed the initial services for this library system, then you already know what these values represent and are ready to begin.

Caution Before installing additional services, ensure that the K2 Admin for initial library services installed on the primary storage server is running. If the service is not running then start the CS Verity K2Admin service from the Administrative Tools >Services panel.

To install additional library system services:

1. On the database server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services are started.
2. Log on as a user who is a member of the local Administrators group on the server where you are installing the additional library system services. (You cannot run Content Services Setup remotely.) This user must satisfy the following requirements:
 - Must not be a member of the Enterprise Admins group
 - Must be a member of the local Administrators group on the server
 - Must be a member of the Domain Admins global group if the server is in a cluster

3. Shut down any applications that use ODBC to connect to the database, such as Enterprise Manager or Query Analyzer.
4. Comment out the following line in **sqlnet.ora** (otherwise you will receive an ambiguous error when you try to start the Storage Manager for the first time):

```
SQLNET.AUTHENTICATION_SERVICES= (NTS)
```

5. If you have not already done so, create an ODBC Data Source Name on the additional storage server (refer to “[Preparing Oracle9i or Oracle 10g Client Connectivity](#)” on page 112 for details.)
6. On the server where you are installing additional library system services (the additional storage server), stop the services for all previously installed library systems.
7. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the machine that will become the initial storage server and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.

Note (FNDTS00071809) If you are performing this installation within a Microsoft Windows Clustering environment, launch the CS Setup program from a command line as follows:

```
setup -log # !logfile.log @ALL
```

where *logfile* is the name of the installation log file you want CS Setup to generate. Due to a bug, launching CS Setup within a clustering environment by double-clicking **Drive:\Setup.exe** fails to generate a log file (**guisetup-username.log**, where *username* is the user running CS Setup).

8. On the Welcome Screen, click Next.
9. On the Software License Agreement screen, click Approve, and then click Next.

10. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. If your database software is Oracle, then click Next. If it isn't, click Cancel to exit Content Services Setup and refer to “Setting Up the Oracle Software” on page 89.
11. On the Select Setup Operation screen, click Install Additional Content Services Components, and then click Next. If you are running CS Setup in a cluster environment, continue at [step 12](#); otherwise, continue at [step 13](#).
12. On the Specify Virtual Server Name screen, enter the virtual server name (VSN), assigned by the domain administrator, that identifies the cluster where you are installing additional services. Then click Next.
13. On the Specify Additional Services Information screen, supply the requested information and click Next:
 - a. Data Source Name: the name you specified when you created an ODBC data source to refer to the database server, which may or may not be the server where you are running CS Setup.
 - b. Tablespace: the name of the tablespace you created as the property repository for this library system.
14. On the Specify Verity Settings screen, note that there is no check box allowing you to choose a Verity K2 Master Admin server. The fully qualified server name and port number you see on this screen are for the Master Admin that governs the Content Search Manager component of the initial services for this library system. This must be the same Master Admin that will govern the Content Search Manager component of the additional services you are installing.

Specify a locale (from the drop-down list) to apply to the Verity collections (indexes) that will be created on this server for this library and click Next. This locale need not be the same as the locale you chose when installing the initial services or other additional services (on other servers) for this library. And the locale can differ from the locale you choose for initial or additional services of other library systems on this server.

15. In the Specify Install Locations screen, supply the requested information and click Next:
 - a. CS Shared Files: the directory that will contain CS shared files. The suggested location is **InstallDrive:\Program Files\FileNet\Shared**. You can choose the location only if the shared files have not been previously installed on the server. Choose a location with at least 20 MB of free space.

Caution Do not specify a network-shared directory. If network connectivity is lost, Content Services components that depend on the shared files will fail.

For each of the following CS components, choose a partition from the drop-down list. The list will show only partitions having at least 200 MB of free space, the amount required for the CS installer.

- b. General Services: the partition where miscellaneous CS utilities will be installed.
- c. Storage Manager: the partition where the additional Storage Manager will be installed.
- d. Content Search Manager: the partition where the additional Content Search Manager will be installed.

Note For optimum performance, install the Content Search Manager in its own partition.

Note For General Services, choose a partition with at least 2.2 GB of free space.

After you click Next, CS Setup will prompt you to create the Shared directory if it hasn't already been created by a previous execution of CS Setup on this server. CS Setup will also prompt you to create a directory (named **filenet**), if it does not already exist, on each partition you selected for General Services, Storage Manager, and Content Search Manager. Click Yes for each prompt.

16. In the Specify Final Setup Values screen, supply the requested information and click Next.
 - a. Installer Name: the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
 - b. Installation Comment: the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is “Content Services 5.5”.
 - c. Select or clear the check box that determines if you want the additional services for this library system to automatically start whenever the server is rebooted.
 - d. Select or clear the check box that determines if the additional services for this library system will automatically start when CS Setup exits. (This check box is available only when you have already installed the initial or additional services of at least one other library system on this server.)
17. The Review Setup Entries screen shows your installation selections in the preceding steps. (You can also see your installation selections in the **setup.inf** file, located at *InstallDrive:\filenet*.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and installing software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

18. In the Installing Components screen, review the results of the installation and click Next. If you are installing the first set of library services on this server, continue at step 20. If you have previously installed initial or additional library system services on this server, continue at step 18.
19. On the Setup Complete screen, select or clear the check box that determines whether you want to perform another setup task within the current execution of CS Setup. Then click Next. If you selected the check box, continue at step 19; otherwise, continue at step 20.
20. On the Select Setup Operation screen, click the option corresponding to the installation task you want to do and click Next:
 - a. To install the initial services for another library system on this server, continue at step 9 of “[Installing Initial Library System Services](#)” on page 123.
 - b. To install additional services on this server for a library system whose initial components are on another server, continue at step 11 (if the server is in a cluster environment) or step 12 (if the server is not in a cluster environment) of this procedure.
 - c. The Reinstall Executables option is generally undertaken only as part of a recovery or move operation, such as moving the database or a set of library services to another server.

- d. To install another display language for a library on this server, continue at step 11 of “Adding a Display Language” on page 134.
21. The appearance of the Exit Setup screen depends on whether you previously installed initial or additional library services on the server prior to the current execution of CS Setup:
- a. If you previously installed library services on the server, simply click Finish to exit CS Setup.
 - b. If you are installing the first set of initial library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

If you installed the additional library services in a cluster environment, continue at “Configuring a Library System Within a Cluster” on page 127.

In a non-cluster environment, after CS Setup exits (and following a restart, if you so specified) perform the steps in “Starting and Monitoring Library Services” on page 128 to start the library services and to enable remote monitoring of them.

View the **cssetup.log** file, located at *InstallDrive:\filenet* by default, to see a log of the individual actions (file copying/deleting, component installation, custom property creation, etc.) in the installation of the additional library services.

View the file **GUISetup-username.log** file, located at *%TEMP%* (where *%TEMP%* is an environment variable for *username*, the user installing Content Services on this server) to see a log of the screen-related actions of CS Setup in the installation of additional library services.

To determine whether the Verity components installed with the additional library services are correctly configured, and to repair the configuration if necessary, run the utility CSK2CFG. Refer to “CSK2CFG” on page 386 for more information.

Adding a Display Language

Content Services can display property names, field names, and system messages in languages other than the default English. Currently, Content Services ships with support files for Chinese, English, French, German, Japanese, Korean, and Spanish (**chinese.nls**, **english.nls**, **french.nls**, **german.nls**, **japanese.nls**, **korean.nls**, and **spanish.nls**). Other language files can be created with direct FileNet assistance. If your company or the vendor for your library system client has created the necessary data file for a particular language and has had it certified by FileNet, you will be able to configure it into your library system so that your users can view Content Services using a language-specific set of display names and text. You can add any number of language files to a library system, thereby supporting concurrent, multiple display languages.

Display Language Installation

Assuming you have the required language data file for a particular language, you can add support for that language by one of two methods:

- Run the CS Setup program as shown in the following procedure within the Content Services Setup program. You should have a screen resolution of 1024x768 on the display where you are running CS Setup. Otherwise, you may have to manually resize one or more CS Setup screens that do not correctly display all their features.
- Run the utility NLSADMIN (refer to “NLSADMIN” on page 438).

To install a display language for an existing library system:

1. Log on to the initial (or an additional) storage server as a user who is a member of the local Administrators group on the server where you are installing the initial or additional library system services. (You cannot run Content Services Setup remotely.) This user must satisfy the following requirements:
 - Must not be a member of the Enterprise Admins group
 - Must be a member of the local Administrators group on the server
 - Must be a member of the Domain Admins global group if the server is in a cluster
2. On the database server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services are started.

Note (FNDTS00071809) If you are performing this installation within a Microsoft Windows Clustering environment, launch the CS Setup program from a command line as follows:

```
setup -log # !logfilename.log @ALL
```

where *logfilename* is the name of the installation log file you want CS Setup to generate. Due to a bug, launching CS Setup within a clustering environment by double-clicking **Drive:\Setup.exe** fails to generate a log file (**guisetup-username.log**, where *username* is the user running CS Setup).

3. Shut down any applications that use ODBC to connect to the database, such as Enterprise Manager or Query Analyzer.
4. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the machine that will become the initial storage server and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.
5. On the Welcome Screen, click Next.
6. On the Software License Agreement screen, click Approve, and then click Next.
7. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. Click Next.
8. On the Select Setup Operation screen, click Install Additional Display Language and then click Next. If you are running CS Setup in a cluster environment, continue at [step 9](#); otherwise, skip to [step 10](#).
9. On the Specify Virtual Server Name screen, enter the virtual server name (VSN), assigned by the domain administrator, that identifies the cluster where you are installing an additional display language. Then click Next.
10. On the Specify Additional Display Language screen, select (highlight) the library system to which you want to add a display language, select a language from the drop-down list, and then click Next.

11. On the Review Setup Entries screen, verify that the Display Language entry shows the library system and display language you selected and click Next.
12. On the Installing Components screen, verify that the display language has been installed and click Next.
13. On the Setup Complete screen, select or clear the check box that determines whether you want to perform another installation task within the current execution of CS Setup. Then click Next. If you cleared the check box, continue at [step 14](#); if you selected the check box, CS Setup continues at [step 15](#).
14. On the Restart Computer screen click Finish to exit CS Setup.
15. On the Select Setup Operation screen, click the option corresponding to the installation task you want to do and click Next:
 - a. To install the initial services for another library system on this server, continue at “[Installing Initial Library System Services](#)” on page 123.
 - b. To install additional services on this server for a library system whose initial components are on another server, continue “[Installing Additional Library System Services](#)” on page 130.
 - c. The Reinstall Executables option is generally undertaken only as part of a recovery or move operation, such as moving the database or a set of library services to another server.
 - d. To install another display language for a library on this server, continue at [step 10](#) of this procedure.

Choosing an Installed Display Language

After you have installed the data files for one or more additional languages, you can set up your site configuration so that a particular language is your system default.

A Display Language property is available in the System, Group, and User objects. With CS Explorer, you can access the System properties and see the value for the Display Language property. If you enable support for other languages, you can select the Display Language property and choose a different language as the system default for your site. You can also specify a language default for particular groups and users.

When a user opens a session from any interface to CS, the system checks the associated User properties, active Group properties, and System properties in this order and uses the first Display Language property value it finds to determine the language in which to display the property names. An Admin user can change the display language by changing the value in the User properties (or by removing that value altogether and accepting the value from the Group properties or System properties).

If no language is specified in the User, Group, or System properties, the system defaults to English. Likewise, if a language data file contains no entry for a particular property name, the name will be displayed in English.

Note Only property names (not property values) are translated. Exceptions to this rule include Yes/No values and the text that appears in the messages that are sent to the Administrators group object and to any User object.

Installing Content Services Unattended

Once you have completed the “[Prerequisites to Installing Content Services](#)” on page 80, you are ready to install the Content Services software. As an alternative to running the interactive Content Services Setup program, you can run an unattended (silent non-interactive) Content Services installation program to install initial and additional library system services.

To perform an unattended installation, you need to do the following:

1. Prepare a parameter file to provide input to the Content Services installer program (refer to “[Preparing a Parameter File](#)” on page 137).
2. Copy the complete installation media to a local physical drive.
3. Invoke the Content Services installer program, using the parameter file as input to the program (refer to “[The dssetup Command Syntax and Options](#)” on page 138 for command options and syntax).

The Content Services installer program used for unattended installation of initial library services, additional library system services, and display languages is named **dssetup.exe** and is on the first *FileNet Content Services* CD-ROM.

After running **dssetup**, continue at “[What to Do Next](#)” on page 129.

Preparing a Parameter File

The parameter file providing input to **dssetup** can be generated in either of two ways.

- Make a copy of a sample parameter file from the first *FileNet Content Services* CD-ROM and edit it to meet the needs of your organization. Refer to
- Run the (interactive) Content Services Setup program as far as the Review Setup Entries screen. But click Cancel, instead of Next, to exit from CS Setup without installing any software. The parameter file containing your installation selections is **setup.inf**, located at *InstallDrive:\filenet*.

The sample parameter files for unattended installation of initial library services and additional library services are available on the first *FileNet Content Services* CD-ROM, as shown in the following table. The table also shows the sample parameter files for unattended reinstalls (refer to “[Reinstalling CS Executable Files Unattended](#)” on page 364).

| Sample Parameter File | Installation Task |
|-------------------------------|--|
| sample_install.inf | Install initial library services |
| sample_addlservice.inf | Install additional library services |
| sample_reinstall.inf | Reinstall initial or additional library services |

The parameter file you specify for the unattended installation of *initial* library services must contain an encrypted password for the administrative user (refer to the AdministrativePassword entry in “[Parameter File Format and Content](#)” on page 139). You can generate the encrypted password from a plaintext password of your choice by running the PWCRYPT utility, found on the first *FileNet Content Services* CD-ROM, as follows:

1. At a command line, run the PWCRYPT command, supplying a plaintext password of your choice:

```
pwcrypt plaintext-password
```

2. PWCRYPT returns the corresponding encrypted password, delimited by doublequotes. Edit the AdministrativePassword entry by setting its value to the encrypted password, including the delimiting doublequotes, as in this example:

```
AdministrativePassword="+:' ,/1.+=%"
```

Using a Parameter File to Install Components

Once you have prepared a parameter file, as explained above, you are ready to run **dssetup** to perform an installation of initial or additional library services or a display language.

To perform an unattended library system installation:

1. On the server where you are installing library system services, stop the services for all previously installed library systems.
2. At a DOS command prompt invoke the **dssetup** command to perform the installation. For example, if the sample file is named **setup.inf**, run the following command:

```
dssetup -f C:\setup.inf -g NO
```

This command will perform an installation according to the parameters specified in **C:\setup.inf** without prompting you for inputs. Refer to “[Parameter File Format and Content](#)” on page 139 for details on the structure of the parameter file.

3. If this is the first installation of library services on the server, reboot the server when **dssetup** exits.
4. Refer to “[Controlling Library System Services](#)” on page 198 for information about stopping, starting, and checking the status of library services.

The dssetup Command Syntax and Options

The **dssetup** command has the following syntax and command options. The name of the parameter file is here assumed to be **setup.inf**, but you can name it something else if you prefer.

```
dssetup [-s hostname] [-i LibrarySystem] [-f drive:\path\filename] [-g {YES | NO}] [-c] [-ti]
```

- c Optional and for Windows SQL only. If you specify -c, **dssetup** will cause the library services to autostart whenever the server is rebooted.
- s Optional. The name of the server where you plan to install library system services. If you do not provide this value, you will be prompted for it. Refer to “[Server Name Requirements](#)” on page 81.
- i Optional. The library system name. This must be the name you assigned to the library system database. If specified, this name overrides a system name supplied in the **setup.inf** file. If you do not provide a name, either by using this flag or by including the name in your **setup.inf** file, you will be prompted for one. Refer to “[Library System Name Requirements](#)” on page 82.
- f Optional. The path and file name of the parameter file. If you do not specify a value, **dssetup** will use **Drive:\filenet\setup.inf**. Refer to “[Parameter File Format and Content](#)” on page 139.
- g (not supported for restore operations). Use this option for unattended installations. If you specify -g NO, **dssetup** will not display prompts. The default value is -g YES (display prompts). If you do not provide a setup file, prompts will be displayed regardless of this option.
- ti Optional. Use this option to cause the CS installer (**dssetup**) to print to the screen each step it is performing.

Parameter File Format and Content

The parameter file stores values that you supply to the **dssetup** command during installation. Using this file you can reduce the number of prompts during interactive installation, or entirely eliminate the interactive process and install the library system unattended.

Every parameter file starts with a **[setup]** header line followed by entries for the keywords shown in the following table. Some entries in the table show a default value in the Description column. This is the value suggested by CS Setup during an interactive installation.

| Keyword | Description |
|---------------------|--|
| Procedure | <p>The Procedure keyword indicates what kind of task Dssetup will perform. It must have one of these values:</p> <p>I = Install a new library system U = Update a previous version of a library system A = Add a component to an existing library system R = Reinstall Content Services executable files</p> |
| Component | <p>The Component keyword indicates what components of the system are affected by the procedure. It must have one or more of the following values:</p> <p>P = Property Manager S = Storage Manager T = Content Search Manager L = Native language support R = Replication Services</p> <p>When several components are specified, they must be separated by commas. For example, a typical multivalue entry would be: Components = P,T,R</p> |
| VcatServerName | The name of the server that contains the database engine (or the virtual server name if the database engine is in a cluster environment) |
| VServerName | The server name on which library system components are to be installed (or the virtual server name if the components are being installed in a cluster environment). Refer to “Server Name Requirements” on page 81 . |
| SystemName | The library system name. Refer to “Library System Name Requirements” on page 82 . |
| IndexTableSpaceName | The name of the tablespace that contains indexes in the library system (Oracle-based Content Services only). |
| DBPassword | The password for the library-system user. Refer to “Creating a Database User” on page 115 . |
| CollateSeq | (Future) For building a SQL engine with a different collate sequence. |

| Keyword | Description |
|-------------------------|---|
| RootName | The name of the top-level installation directory for all installed library system components. If you are installing library system components on more than one drive, a directory with this name will be created at the root level of each drive and the selected components will be installed there. |
| InstDrive | The partition containing the top-level installation directory for all installed library system components. |
| SharedDir | The directory containing the Content Services shared files Default directory path is C:\Program Files\FileNet\Shared |
| SystemLanguage | Display language for the installation interface (English only). |
| AdministrativeUser | The name you want to assign to the initial administrative user in the library system. The default user name is <i>Admin</i> . |
| AdministrativePassword | The encrypted password for the administrative user. |
| AdministrativeGroup | The name you want to assign to the initial administrative group in the library system. The default group name is <i>Administrators</i> . |
| UniversalGroup | The name you want to assign to the universal group. The default value is <i>General Users</i> . |
| StorageCategory | The name you want to assign to the initial storage category. The default name is <i>Documents</i> . |
| ArchiveCategory | The name you want to assign to the initial archive category. The default name is <i>Archive</i> . |
| DocClass | The name you want to assign to the default document class. The default name is <i>General</i> . |
| StorageRepositoryDrive | The drive assigned to storage repository files for the library system. |
| IndexRepositoryDrive | The drive assigned to index files for the library system. |
| ArchiveRepositoryDrive | The drive assigned to archive files for the library system. |
| VerityMasterAdminServer | The fully qualified name of the server where the Verity K2 Master Admin that controls the library system is installed. |
| VerityMasterAdminPort | The port number assigned to the Verity K2 Master Admin. The default port assignment is 9550. |
| VerityIndexLanguage | The Verity locale (language) assigned to the Verity collections on the server where the library system is being installed. The default language is <i>English</i> . |
| Installer | The name entered by the person installing. The default value is <i>Administrator</i> . |

| Keyword | Description |
|-----------|--|
| Comment | The comment entered by the person installing. The default value is <i>Content Services 5.5</i> . |
| AutoStart | This keyword controls whether the library system services will be started automatically when the server is rebooted or must be started manually. Valid values are AutoStart = 1 (start services on reboot) and AutoStart = 0 (manual startup). |
| StartNow | This keyword controls whether the library system services will be started automatically as soon as the dssetup completes. Note that this keyword controls the starting behavior of all but the first library system you install on a server (the server must be restarted before the first library system can be used). |
| LogFile | The path to the file that records the actions of dssetup . The default path is C:\filenet\cssetup.log. |
| Verbose | The keyword that controls whether the information recorded during execution of dssetup is detailed (Verbose = 1) or not (Verbose = 0). |

Examples of Parameter Files

The following files are examples of parameter files that you can edit according to the needs of your organization to install Content Services components. (Note that the IndexTableSpaceName entry in the file applies only when you are installing Content Services on an Oracle platform).

Example parameter file to install initial library services:

```
[setup]
Procedure=I
Component=P, S, T, L, R
VcatServerName=DBServer
VServerName=localhost
SystemName=libraryname
IndexTableSpaceName=
DBPassword="Encrypted password string"
RootName=filenet
InstDrive=C:
SharedDir=C:\Program Files\FileNet\Shared
SystemLanguage=English
AdministrativeUser=Admin
AdministrativePassword="Encrypted password string"
AdministrativeGroup=Administrators
UniversalGroup=General Users
StorageCategory=Documents
ArchiveCategory=Archive
DocClass=General
StorageRepositoryDrive=C:
IndexRepositoryDrive=C:
VerityMasterAdminServer=K2AdminServer.domain.local
VerityMasterAdminPort=9550
```

```
VerityIndexLanguage=English
Installer=Administrator
Comment=Content Services 5.5
AutoStart=0
StartNow=0
LogFile=C:\filenet\cssetup.log
Verbose=1
```

Example parameter file to install additional library services:

```
[setup]

Procedure=I
Component=S, T, R
VcatServerName=DBServer
VServerName=localhost
SystemName=libraryname
IndexTableSpaceName=
DBPassword=
RootName=filenet
InstDrive=C:
SharedDir=C:\Program Files\FileNet\Shared
SystemLanguage=
AdministrativeUser=
AdministrativePassword=
AdministrativeGroup=
UniversalGroup=
StorageCategory=
ArchiveCategory=
DocClass=
StorageRepositoryDrive=C:
IndexRepositoryDrive=C:
VerityMasterAdminServer=K2AdminServer.domain.local
VerityMasterAdminPort=9550
VerityIndexLanguage=
Installer=Administrator
Comment=Content Services 5.5
AutoStart=0
StartNow=0
Verbose=1
LogFile=c:\filenet\cssetup.log
```

Upgrading the Oracle Database

Upgrading your database to Oracle 10g 10.2.0.1.0 requires that you upgrade your software and migrate your existing database information to the new version. You need to choose one of these migration methods:

- Export/Import
- Database Upgrade Assistant
- Manual upgrade

For more information on how to migrate your existing data, or information on the Oracle Database Upgrade Assistant, refer to your Oracle documentation.

Note FileNet strongly recommends the Export/Import migration method to migrate your database to Oracle 10g 10.2.0.1.0.

If you move your Oracle database to another server before upgrading from CS 5.4 to CS 5.5, make sure to use the same name and case when restoring the database on the new server. For example, if a tablespace is named DEVELOP on server A, it must also be named DEVELOP on server B. If you do not use the same name and case, the upgrade to Content Services 5.5 will fail.

You can use Oracle's Export/Import feature to export your existing Oracle9i 9.2.0 database to a file, then import the file into Oracle 10g 10.2.0.1.0. One advantage of this method is the ability to move the database from one machine to another, such as when you update your hardware along with the software. Refer to [“Exporting/Importing a Windows Oracle Database” on page 148](#) or [“Exporting/Importing a UNIX Oracle Database” on page 150](#) for more information.

The Oracle Database Upgrade Assistant (DUA) provides a completely automated upgrade of your database, using a graphical user interface to guide you through each step of the process. The DUA runs the appropriate upgrade script for your current release, deletes any obsolete initialization parameters from your initialization parameter file, and (optionally) configures your **listener.ora** file.

Note You can only use the DUA to upgrade a database on the same server. If you are moving the database to a new server, you must use the Export/Import Method.

Manually upgrading is a complex process, beyond the scope of this manual. Refer to your Oracle documentation for complete information.

Upgrading to Oracle 10g 10.2.0.1.0 on Windows

The following procedure provides minimal instructions for upgrading to Oracle 10g 10.2.0.1.0 from Oracle9i 9.2.0 on a Windows platform. Details are given only where Content Services has specific requirements pertaining to the Oracle installation and data migration. Refer to your Oracle 10g 10.2.0.1.0 documentation for complete information on upgrade options and procedures.

Caution Oracle recommends that you use a new Oracle home directory when you upgrade your current 9.2.0 database engine to Oracle 10g 10.2.0.1.0. If you do not use a new Oracle home when upgrading, you must first uninstall your current Oracle9i 9.2.0 release before upgrading. You do not need to uninstall your current Oracle9i 9.2.0.4 if you are installing Oracle9i 9.2.0 to a new Oracle home.

Note Direct upgrade from only Oracle9i 9.0.1 or higher / Oracle9i R2 9.2.0 or higher is supported.

To upgrade the database to Oracle 10g 10.2.0.1.0 from Oracle9i 9.2.0:

1. On the property server, stop all applications that access Oracle 9i 9.2.0.
2. Stop all Oracle 9i 9.2.0 services.
3. Insert the Oracle 10g 10.2.0.1.0 CD-ROM into the CD-ROM drive. The Autorun main menu opens. Click Install/Deinstall Products.
4. On the Welcome screen, select “Advanced Installation” radio button and then click Next.
5. On the Select Installation Type screen, select Custom and then click Next.
6. On the Specify Home Details screen, change the Destination Name to a name other than other than the current (9.2.0) name. Change the Destination Path to another path, such as c:\oracle\ora10g, and then click Next.
7. The Available Product Components screen lists all components available for installation. Select the components you wish to install (see below) and click Next.

Except where noted, the following items are the bare minimum requirements for Content Services:

- Oracle Database 10g 10.2.0.1.0
- Oracle Net Services 10.2.0.1.0
 - Oracle Net Listener 10.2.0.1.0
- Oracle Call Interface (OCI) 10.2.0.1.0
- Oracle Windows Interfaces 10.2.0.1.0
 - Oracle ODBC Driver 10.2.0.1.0 (optional)

Note Although Oracle ODBC Driver is neither certified nor used by Content Services, and need not be installed, you may want to install it anyway since it contains ODBC Test, a useful troubleshooting tool for testing ODBC connectivity.

You may also want to install other components for added flexibility or administration, such as Oracle Partitioning or Enterprise Manager Products.

8. On the Create Database screen, select “Install database software only” and click Next.

9. Read the summary and click on Install.
10. Configure Listener.
 - Launch the Oracle Net Configuration Assistant tool.
 - On Welcome screen, select “Listener Configuration” and click Next.
 - Select “Add”, click Next.
 - Specify “Listener Name”, click Next.
 - The Selected protocols list should contain TCP as one of the values. Click Next.
 - Specify the TCP/IP Port Number. Click Next.
 - Select No for configuring another listener.
 - Click Finish.

Caution

- After upgrading and migrating to Oracle Database 10g Release 2, the CONNECT role will only have the CREATE SESSION privilege; the other privileges granted to the CONNECT role in earlier releases will be revoked during the upgrade. After migrating database using export/import method, log on to new database server as SYS user with SYSDBA privileges and execute the following SQL script once per each library system. Before running the script, substitute your library system (tablespace) name in place of LIBRARYNAME.

```
grant alter session to LIBRARYNAME;  
grant create cluster to LIBRARYNAME;  
grant create database link to LIBRARYNAME;  
grant create sequence to LIBRARYNAME;  
grant create synonym to LIBRARYNAME;  
grant create table to LIBRARYNAME;  
grant create view to LIBRARYNAME;
```

- If Oracle database is migrated using Export/Import method, then CS executables must be re-installed for each library system.

Upgrading to Oracle 10g 10.2.0.1.0 on UNIX

The following procedure provides minimal instructions for upgrading to Oracle 10g 10.2.0.1.0 from Oracle9i 9.2.0 as a part of library system upgrade. Details are given only where Content Services has specific requirements pertaining to the Oracle installation and data migration. Refer to your Oracle 10g 10.2.0.1.0 documentation for complete information on upgrade options and procedures.

Caution Oracle recommends that you use a new Oracle home directory when you upgrade your current Oracle9i 9.2.0.4 database engine to Oracle 10g 10.2.0.1.0. If you do not use a new Oracle home when upgrading, you must first uninstall your current Oracle9i 9.2.0. release before upgrading. You do not need to uninstall your current Oracle9i 9.2.0 if you are installing Oracle9i 9.2.0 to a new Oracle home.

Note Direct upgrade from only Oracle 9i 9.0.1.4 or higher / Oracle 9i R2 9.2.0 or higher is supported.

To upgrade the database to Oracle 10g 10.2.0.1.0 from Oracle9i 9.2.0:

1. On the property server, stop all applications that access Oracle 9i 9.2.0.
2. Stop all Oracle 9i 9.2.0 processes.
3. Log on to the property server as user oracle.
4. Access the mount point directory for the device containing the Oracle 10g CD. For example:


```
cd /cdrom
```
5. Run the Oracle Universal Installer:


```
./runInstaller
```
6. On the Welcome screen, select “Advanced Installation” radio button and then click Next.
7. On the Select Installation Type screen, select Custom and then click Next.
8. On the Specify Home Details screen, change the Destination Name to a name other than the current (9.2.0) name. Change the Destination Path to another path, such as **/oracle/product/10.2.0**, and then click Next.

The Available Product Components screen lists all components available for installation. Select the check boxes of the following components for installation. Except where noted, these are the bare minimum requirements for Content Services:

- Oracle Database 10g 10.2.0.1.0
- Oracle Net Services 10.2.0.1.0
 - Oracle Net Listener 10.2.0.1.0
- Oracle Call Interface (OCI) 10.2.0.1.0
- Oracle ODBC Driver 10.2.0.1.0 (optional)

Note Although Oracle ODBC Driver is neither certified nor used by Content Services, and need not be installed, you may want to install it anyway since it contains ODBC Test, a useful troubleshooting tool for testing ODBC connectivity.

You may also want to install other components for added flexibility or administration, such as Advanced Replication, Oracle Partitioning, or Enterprise Manager Products.

9. On Product Specific Prerequisite Checks screen, click Next.
10. On Privileged Operating System Groups screen, specify the OSDBA and OSOPER groups and click Next.
11. On the Create Database screen, select “Install database software only” and click Next.
12. Read the summary and click on Install.
13. Configure Listener.
 - Launch the Oracle Net Configuration Assistant tool.
 - On Welcome screen, select “Listener Configuration” and click Next.
 - Select “Add”, click Next.
 - Specify “Listener Name”, click Next.
 - The Selected protocols list should contain TCP as one of the values. Click Next.
 - Specify the TCP/IP Port Number. Click Next.
 - Select No for configuring another listener.
 - Click Finish.

Caution

- After upgrading and migrating to Oracle Database 10g Release 2, the CONNECT role will only have the CREATE SESSION privilege; the other privileges granted to the CONNECT role in earlier releases will be revoked during the upgrade. After migrating database using export/import method, log on to new database server as SYS user with SYSDBA privileges and execute the following SQL script once per each library system. Before running the script, substitute your library system (tablespace) name in place of LIBRARYNAME.

```
grant alter session to LIBRARYNAME;
grant create cluster to LIBRARYNAME;
grant create database link to LIBRARYNAME;
grant create sequence to LIBRARYNAME;
grant create synonym to LIBRARYNAME;
grant create table to LIBRARYNAME;
grant create view to LIBRARYNAME;
```

- If Oracle database is migrated using Export/Import method, then CS executables must be re-installed for each library system.

Exporting/Importing a Windows Oracle Database

Use this method to migrate an existing Oracle database on a Windows server to a new version of Oracle, or to move an existing database to another server running a Windows version of Oracle. First export the existing database information, then import the database to a another Oracle installation on the same server or a different server.

Caution

- If you are migrating the database to another server, FileNet strongly recommends that you retain the old server name, especially for library systems participating in Replication Services or library systems with Panagon stored searches and search templates. If you need to move your database to a server with a different name, refer to [“Moving the Database” on page 352](#).
- This procedure requires you to use the Windows Registry Editor, something only experienced system administrators should do. Making mistakes while editing the registry can make your system unstable or worse!

Perform the following steps on the existing database server.

To export your database:

1. Log on to the database server as user *oracle*.
2. Create a file named **export.par** that contains the following lines:

```
FILE=export_path\export_file.dmp
LOG
OWNER=library_system_name
RECORD=Y
```

3. Create a file name **import.par** that contains the following lines:

```
FILE=export_path\export_file.dmp
FROMUSER=library_system_name
TOUSER=library_system_name
```

4. Launch the Export utility as follows:

```
exp system/system_password PARFILE=drive:\path\export.par
```

If you are upgrading the database, perform the following steps on the existing server. If you are moving the database, perform these steps on the new database server.

To import your database:

If you are upgrading the database, perform these steps on the existing server. If you are moving the database, perform these steps on the new database server:

1. Install (or upgrade to) Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0. Refer to [“Installing Oracle9i on Windows” on page 92](#) or [“Installing Oracle 10g on Windows” on page 97](#).
2. Create a database. Make sure that you use the database instance name (SID) and Global Database Name (Service Name) from your exported database. Refer to [“Creating an Oracle9i Database” on page 102](#) or [“Creating an Oracle 10g Database” on page 104](#).

3. Create a tablespace. Make sure you use the same name and case as in the exported database. For example, if the tablespace is named DEVELOP on server A; it must also be named DEVELOP on server B. Refer to “Creating a Property Repository Tablespace” on page 106 for instructions.
4. Create a database user. Refer to “Creating a Database User” on page 115.
5. Import the user's data by running the following SQL statement:

```
imp library_system_name/library_system_name  
PARFILE=full_path_to_param_file\import.par
```

Note Importing the database may cause some of your stored procedures to be marked invalid. To restore their validity, run the script **ORACLE_HOME\rdbms\utlrp.sql** or run the Content Services utility **CAT_PROC** after you have installed Content Services 5.5.

6. If you are moving the database to a new server, import the FileNet registry key on the new database server. Using REGEDIT, select:

```
HKEY_LOCAL_MACHINE > SOFTWARE
```

From the Registry menu, select Import Registry File. Use the saved registry file you exported earlier.

7. Configure Oracle client connectivity to the server where you imported the database. Refer to “Connecting to the Oracle Database” on page 109.

Exporting/Importing a UNIX Oracle Database

Use this method to migrate an existing Oracle database on a UNIX server to a new version of Oracle, or to move an existing database to another server running a UNIX version of Oracle. First export the existing database information, then import the database to a another Oracle installation on the same server or a different server.

Caution If you are migrating the database to another server, FileNet strongly recommends that you retain the old server name, especially for library systems participating in Replication Services or library systems with Panagon stored searches and search templates. If you need to move your database to a server with a different name, refer to [“Moving the Database” on page 352](#).

Perform the following steps on the existing database server.

To export your database:

1. Log on to the database server as user *oracle*.
2. Create a file named **export.par** that contains the following lines:

```
FILE=export_path/export_file.dmp
LOG
OWNER= library_system_name
RECORD=Y
```

3. Create a file name **import.par** that contains the following lines:

```
FILE=export_path/export_file.dmp
FROMUSER=library_system_name
TOUSER=library_system_name
```

4. Run the export utility. Enter the following command:

```
exp system/system_password PARFILE=device:/path/export.par
```

If you are upgrading the database, perform these steps on the existing server. If you are moving the database, perform these steps on the new database server.

To import your database:

1. Install (or upgrade to) Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0. Refer to [“Installing Oracle9i on UNIX” on page 94](#) or [“Installing Oracle 10g on UNIX” on page 99](#).
2. Create a database. Make sure that you use the database instance name (SID) and Global Database Name (Service Name) from your exported database. Refer to [“Creating an Oracle9i Database” on page 102](#) or [“Creating an Oracle 10g Database” on page 104](#).
3. Create a tablespace. Make sure you use the same name and case as in the exported database. For example, if the tablespace is named DEVELOP on server A; it must also be named DEVELOP on server B. Refer to [“Creating a Property Repository Tablespace” on page 106](#) for instructions.
4. Create a database user. Enter the following statement in SQL*Plus:

```
create user library_system_name identified by library_system_name
default tablespace library_system_name
```

```
temporary tablespace temp_tablespace;  
grant connect, dba to library_system_name;
```

Note After the import, you can revoke dba privileges from the CS user, or you can let the CS 5.4 upgrade software do this automatically for you.

5. Import the user's data by running the following SQL statement:

```
imp library_system_name/library_system_name PARFILE=full_path_to_param_file/  
import.par
```

Note Importing the database may cause some of your stored procedures to be marked invalid. To restore their validity, run the script **\$ORACLE_HOME/rdbms/utlrp.sql** or run the Content Services utility **CAT_PROC**.

6. Configure Oracle client connectivity to the server where you imported the database. Refer to “Connecting to the Oracle Database” on page 109.

Upgrading Content Services

Perform the upgrade tasks explained in the following topics, in the order shown:

1. Refer to “[Upgrade Overview](#)” on page 152 and “[Upgrade Constraints](#)” on page 153, for information on what an upgrade involves.
2. Refer to the “[Upgrade Roadmap](#)” on page 75 for a complete checklist of the tasks required when upgrading a library system.
3. Refer to “[Preparing for the Upgrade](#)” on page 154 for a detailed list of the tasks you need to perform on the server where you are going to upgrade Content Services.
4. Refer to “[Upgrading Library System Services](#)” on page 157 if you are going to upgrade library services using the Content Services Setup program.
5. Refer to “[Rebuilding Content Search Repositories](#)” on page 160 for instructions on rebuilding content search repositories after all library systems have been upgraded.
6. Refer to “[Upgrading Content Services Unattended](#)” on page 160 if you are going to upgrade initial or additional library system services using a non-interactive installation program.

Upgrade Overview

The upgrade path to Content Services 5.5 is from Content Services 5.4 only. If your library system is an earlier version than 5.4 you must first upgrade to 5.4 according to the instructions in the CS 5.4 documentation, then upgrade to Content Services 5.5. Content search repositories must be rebuilt after upgrading to CS 5.5.

If you want to upgrade the database to Microsoft SQL Server 2005 or Oracle10g 10.2.0.1.0, then this must be done after upgrading to CS 5.5

If you have Content Services 5.4 installed in a non-cluster environment and want to upgrade it into a cluster environment, you must first upgrade to Content Services 5.5 *outside* the cluster, and then move Content Services into the cluster environment.

Upgrading Content Services means upgrading the initial and additional sets of services for each of your library systems. To upgrade a set of library services, you run the Content Services Setup program (or run the CS installation program noninteractively via a parameter file). When you upgrade a set of services for a library system, CS Setup automatically installs/upgrades all CS 5.4 library services components (Storage Manager, Content Search Manager, and Replication Services) of that library system on the server where it is running.

Regardless of how you perform the upgrade, make sure that the server where you are upgrading to CS 5.5 is running Windows 2000 or Windows 2003, not Windows NT. That is, you must upgrade the operating system, as needed, before upgrading to CS 5.5. For best results, you should move you set of library services to a clean install of Windows 2000 or Windows 2003 rather than just upgrading the operating system and required software. You can choose to perform a full Windows 2000 or Windows 2003 install on the same server, or you can move the library services to a new server running Windows 2000 or Windows 2003.

If you want to upgrade a set of library services from CS 5.4 on one server to CS 5.5 on a new server, it is recommended that you first move the CS 5.4 library services from the old server to the new server according to the

instructions in the CS 5.4 documentation. Then upgrade the library services from CS 5.4 to CS 5.5 on the new server.

A library system upgrade is not complete until you have upgraded every set of services for that library system. While you are upgrading the library services, the library system is inaccessible to users. During the time you are upgrading the services for that library.

It is recommended that you upgrade all the initial and all additional services for one library system before you upgrade the initial or additional services for another library system.

If initial or additional services for more than one CS 5.4 library system are installed on a server, you must upgrade all the services to CS 5.5 before using any of the library systems on the server. That is, initial or additional component of every library system on a server must be of the same version.

A consequence of this requirement is that once you have run the first CS 5.4-to-CS 5.5 upgrade of initial or additional library services for a given library on server A, the indexes managed by a CS 5.4 Content Search Manager for the same library on some other server B will be unavailable until all sets of services for all libraries on server B have been upgraded from CS 5.4 to CS 5.5.

Moreover, you cannot have library systems of different versions (such as Content Services 5.4 and Content Services 5.5) replicating documents among themselves. Therefore, if you are upgrading a library system that participates in replication (either inside or outside an administrative domain), you must upgrade all the other library systems, either on the same server or on other servers, before you can resume replication.

Upgrade Constraints

The order in which you upgrade your servers is constrained by the location of the Verity K2 Master Admin. The server upon which the K2 Master Admin resides must be upgraded first, then the other non-master servers can be upgraded in any order.

To prepare Verity K2 for the upgrade, you must run a program, `prepk2upg`, to detach and deregister all K2Brokers, K2Servers, and collections.

Preparing for the Upgrade

Note At least **4 weeks** prior to the upgrade, the TSP (Technical Service Provider) **must** schedule the upgrade with the FileNet Upgrade/Install Assurance Team (U/I Team). To access the U/I Team's latest scheduling procedures go to [FileNet Worldwide Customer Support](#) web site.

Note: FileNet Content Services Hotfix Pack CS-5.4.0-002 is required while upgrading from FileNet CS 5.4 to CS 5.5.

Before you can upgrade Content Services, you need to perform the following steps on each server where you will be upgrading initial or additional library services.

Before you can upgrade Content Services, you need to perform the following steps on each server where you will be upgrading initial or additional library services.

Read the Content Services Documentation:

1. Review the hardware and software requirements for the new version of Content Services. Refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#).

Back up the servers and library systems:

1. Notify users to check in all checked-out items and log off.
2. Use CS Admin Tools to ensure that the Replication Enabled? property (on the Properties tab of the System object) is set to Yes for each library system that participates in replication, and No for each library system that does not.
3. Use CS Admin Tools to put the library system in single-user mode. A library system is in single-user mode when:
 - The Administrators group is your active group
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
4. Log out of CS Admin Tools.
5. In a non-cluster environment, use the **dsstop** utility or the Content Services Manager applet (refer to your *Content Services 5.4 Administrator's Guide* for details) to shut down all the CS 5.4 library system services on each server. In a cluster environment, use Microsoft Cluster Administrator to shut down the services.
6. Start SQL*Plus and connect to your database instance when prompted for a username:


```
sys/sys_pwd@service_name as sysdba
```
7. To ensure that all sessions are closed, use your database tools to run the following SQL statements:


```
select * from librarysystem.sessiontab;
```

 If there are entries in the results, run:


```
delete * from librarysystem.sessiontab;
```
8. (Strongly recommended but not required) Uninstall the CS client libraries on each machine accessing CS 5.4 before upgrading to CS 5.5.

Note This must be done in case where IDM Desktop and CS Admin tool both are co-located. After the upgrade is done, run the CS 5.5 client libraries installer to install the client libraries.

9. Perform a full offline backup of the database, making sure to back up the files listed in the **v8files.log** spool file that you generated above.

Caution If you encounter any problems while upgrading Content Services, you will need to restore the database from this backup. Therefore, make sure you back up your database now as a precaution. See the *Oracle9i Backup and Recovery Guide* for more information.

Satisfy the hardware and software requirements for the Content Services 5.5 servers:

1. Be sure that the new servers satisfy the hardware and software requirements for Content Services listed in [CS 5.5.0 Windows Oracle Platform Requirements](#).
2. Make sure the Windows operating system is on the C: drive of the server where you will upgrade to a CS 5.5 set of additional library services whose Content Search Manager component will be governed by a Verity K2 Master Admin on some other server. Otherwise, the upgrade may fail.
3. Start SQL*Plus and connect to your database instance when prompted for a username:

```
sys/sys_pwd@service_name as sysdba
```

4. Run the following SQL statement to add 20% (or at least 20 MB) more free space to your SYSTEM tablespace and to the tablespaces where you store rollback segments, if necessary.

Note Upgrading to a new release requires more space in your SYSTEM tablespace and in the tablespaces where you store rollback segments. If you have enough space on your system, then consider adding more space to these tablespaces. In general, you need at least 20MB of free space in your SYSTEM tablespace to upgrade. If you run out of space during the upgrade, then you will need to repeat the upgrade with adequate space.

If your Oracle database server is running Windows:

```
ALTER TABLESPACE system ADD DATAFILE 'drive:\path_to_system_dat\system.dat' SIZE
[current_size + 20%] REUSE
AUTOEXTEND ON NEXT 5m MAXSIZE UNLIMITED;
```

If your Oracle database server is running UNIX:

```
ALTER TABLESPACE system
ADD DATAFILE '/path_to_system_dat/system.dat' SIZE [current_size + 20%] REUSE
AUTOEXTEND ON NEXT 5m MAXSIZE UNLIMITED;
```

5. Determine the files that you will back up by executing the following SQL statements:

```
SPOOL v8files.log;
SELECT member FROM v$logfile;
SELECT name FROM v$datafile;
SELECT value FROM v$parameter WHERE name='control_files';
SPOOL OFF;
```

Optionally, you can create index tablespaces to hold the indexes that will be built in your upgraded CS 5.5 library systems, one index tablespace per library system. If you want index tablespaces, continue at [step 6](#); otherwise skip to [step 9](#).

6. When creating an index tablespace, the data file for the tablespace must be at least 1 MB if you specify AUTOALLOCATE, or at least 15 MB if you don't specify AUTOALLOCATE. For example, to create a locally managed Oracle9i 9.2.0 80 MB index tablespace **lib1x**, enter the following SQL command.

- For Oracle running on Windows:

```
create tablespace lib1x
datafile 'install_drive:\[path_to_tablespace_files]\lib1x.dbf' size 80m
autoallocate;
```

- For Oracle running on UNIX:

```
create tablespace lib1x
datafile '[mount_point]/lib1x.dbf' size 80m
autoallocate;
```

7. Read and edit the SQL script **indextbsp.sql**, located on the first *FileNet Content Services* CD-ROM, for each library system to be upgraded. Then, as the SYS user, run **indextbsp.sql**. The output is another script, **alterIndexTbsp.sql**.

8. As the SYS user, run **alterIndexTbsp.sql**.

9. Run the SQL statement SHUTDOWN IMMEDIATE on the database.

10. Perform a full offline backup of the database, making sure to back up the files listed in the **v8files.log** spool file that you generated above.

Caution If you encounter any problems while upgrading Content Services, you will need to restore the database from this backup. Therefore, make sure you back up your database now as a precaution. See the *Oracle8i Backup and Recovery Guide* for more information.

11. If you haven't already done so, add the database server and all initial and additional storage servers to a Windows 2000 or Windows 2003 domain. To add a server to a domain, do the following:

- a. Open the Windows 2000 or Windows 2003 control Panel and start the System tool.
- b. Go to the Network Identification tab and click Properties.
- c. Select the Domain option button and type your domain name. For example, to join the MYDOM.LOCAL domain, type MYDOM and click OK.

12. Reboot the server.

Preparing Verity K2 for the Upgrade

To prepare Verity K2 for this upgrade, you must run the program `prepk2upg` once before the upgrade process is started. This program will identify the K2 master admin server and will detach and deregister all the K2 components (brokers, servers, and collections) on that K2 master admin for all the library systems on all the servers.

Usage

prepk2upg [-l]

- | | |
|----------------------------|---|
| prepk2upg -l | lists the master admin, lists all of the admin servers, collections and k2servers. |
| prepk2upg (without the -l) | detaches and deregisters all the K2 components on all servers, and lists the master admin server. |

To verify that Verity K2 has been prepared for the upgrade, run `prepk2upg -l` and only the admin servers will be displayed.

Upgrading Library System Services

If you have not already performed the steps in “[Preparing for the Upgrade](#)” on page 154, do so now. Also, if you previously ran CS Setup for the first time on this server to upgrade library services, restart the server now to complete the upgrade before running CS Setup again. Then perform the steps in the following procedure to upgrade a set of library services on this server from CS 5.4 to CS 5.5.

Warning CS Setup will fail if you did not upgrade the K2 master admin server first.

Warning `prepk2upg` should be run before upgrading the K2 master admin server.

To upgrade library services from CS 5.4 to CS 5.5:

1. On the database server, verify that Oracle is running:
 - a. If Oracle is installed on a Windows server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services is started.
 - b. If Oracle is installed on a UNIX server, see that the Oracle processes are running by the following command:

```
ps -eaf | grep ora
```

2. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the initial storage server and double-click **Drive:\Setup.exe** to launch the CS Setup program.

Note (FN DTS00071809) If you are performing this installation within a Microsoft Windows Clustering environment, launch the CS Setup program from a command line as follows:

```
setup -log # !logfile.log @ALL
```

where *logfile* is the name of the installation log file you want CS Setup to generate. Due to a bug, launching CS Setup within a clustering environment by double-clicking **Drive:\Setup.exe** fails to generate a log file (**guisetup-username.log**, where *username* is the user running CS Setup).

3. On the Welcome screen, click Next.
4. On the Software License Agreement screen, click Approve, and then click Next.
5. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. If your database software is Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0,

then click Next. If it isn't, click Cancel to exit Content Services Setup and refer to “Setting Up the Oracle Software” on page 89.

6. On the Select Setup Operation screen, click Upgrade Content Services Components and then click Next.
7. On the Specify Upgrade Information screen, select the library system you want to upgrade from the Tablespace Name list of libraries whose components on this server have not yet been upgraded. (You can only select one library at a time.) If you have modified the ODBC data source name, net service name, and/or have exported the database to a new server (creating a database user with a temporary password), select the check box and enter the current password for the library system user. Then click Next.
8. On the Specify Verity Settings screen, specify the Verity K2 Master Admin server name and port (as needed) and the Verity locale (language), as described below, and then click Next.
 - a. If the screen has no check box, skip to step b, since the library services you are installing will be governed by an existing Master Admin (whose server name and port are shown on the screen) either on this server or some other server. If the screen does have a check box, do one of the following:
 - Select the check box and specify the fully qualified server name and port of an existing Master Admin on some other server to govern the library services you are upgrading.
 - Clear the check box to install a Master Admin (the server name and port are shown on the screen) with the library services you are upgrading on the server.
 - If you are upgrading library services in a cluster, clear the check box and type the fully qualified cluster name (*cs_clustername_verity.domain*) and port.
 - b. The locale you choose (from the drop-down list) applies only to the Verity collections (indexes) on this server for this library system, and can differ from the locale for services of this library system on other servers. It can also differ from the locale for services of other library systems on this server.
9. The Review Setup Entries screen shows your upgrade selections in the preceding steps. You can also see your installation selections in the **setup.inf** file, whose location is shown in the last line of the Review Setup Entries screen. If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and upgrading software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

10. On the Installing Components screen, review the results of the upgrade and click Next. If you are performing the first upgrade of library services on the server, skip to [step 12](#); otherwise, continue at [step 11](#).

11. On the Setup Complete screen, select or clear the check box that determines whether you want to perform another CS setup task within the current execution of CS Setup. Then click Next. If you selected the check box, skip to [step 13](#); if you cleared the check box, continue at [step 12](#).
12. The appearance of the Exit Setup screen depends on whether you previously upgraded library services on the server prior to the current execution of CS Setup:
 - If you previously upgraded library services on the server, simply click Finish to exit CS Setup.
 - If you have not previously upgraded library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

13. On the Select Setup Operation screen, the options available to you depend on the current upgrade status of the various sets of library services installed on the server, as noted below. Click the option corresponding to the installation task you want to perform and then click Next:

The following option is the only one available if at least one set of library services on the server has not yet been upgraded from CS 5.4 to CS 5.5:

- Upgrade Content Services Components: Click this option and continue at [step 7](#) to upgrade another set of library services on this server.

If all sets of library services on the server have been upgraded from CS 5.4 to CS 5.5, the following options are available:

- Install Initial Content Services Components: Click this option and continue at [“Installing Initial Library System Services” on page 123](#) to install the initial services for a new library system on this server.
- Install Additional Content services Components: Click this option and continue at [“Installing Additional Library System Services” on page 130](#) to install additional services on this server for a library system whose initial services are on another server.
- Reinstall Executables: Click this option only as part of a recovery or move operation, such as moving the database or a set of library services to another server.
- Install Additional Display Language: Click this option and continue at [“Adding a Display Language” on page 134](#) to install another display language for a set of library services on this server.

14. If this is the first upgrade of library services on the server, reboot the server when CS Setup exits.

View the **cssetup.log** file, located at *InstallDrive:\filenet* by default, to see a log of the individual actions (file copying/deleting, component installation, custom property creation, etc.) in the upgrade of library services.

View the file **GUISetup-username.log** file, located at **%TEMP%** (where **%TEMP%** is an environment variable for *username*, the user upgrading Content Services on this server) to see a log of the screen-related actions of CS Setup in the upgrade of library services.

You must only run CSK2CFG and re-attach all collections for all libraries, and then stop and restart CS services. Refer to [“CSK2CFG” on page 386](#) for more information. You may need to stop and start CS services twice for the

changes to take effect. You can then check to make sure everything is attached with CSK2CFG, or `rcadmin>hierarchyview>all`.

After upgrading every set of services for each library system on all the servers, continue at [“Rebuilding Content Search Repositories”](#) on page 160.

If you are doing chain upgrade (from a pre-CS 5.4 installation), after upgrading every set of services for each library system on all the servers, continue at [“Rebuilding Content Search Repositories”](#) on page 160.

Rebuilding Content Search Repositories

Before you can perform content searches, you must rebuild the content search repositories (Verity collections) in the library systems you have upgraded. Be sure you have upgraded each set of services for every library system *before* you rebuild any content search repositories.

To rebuild content search repositories:

1. For each upgraded library system, find all the servers where the library system has content search repositories. To find the servers, log onto the library system via CS Admin Tools, select each repository number (“1”, “2”, etc.) under the Content Search Repositories object in the tree view, and note the corresponding Content Search Repository Name (**index001**, **index002**, etc.) and Server entries in the list view.
2. For each upgraded library system, run the CSADMIN utility (see [“CSADMIN”](#) on page 384 for details) on the server where **index001** is located. Invoke the utility as follows:

```
csadmin propertyserver libraryname -r
```

where *propertyserver* is the property server name and *libraryname* is the name of the library system. Then respond 'y' at each CSADMIN prompt to rebuild all the content search repositories of the library system on the server.

Caution For each library system, you must rebuild **index001** first. Otherwise, content searches in the upgraded library systems will fail.

3. For each upgraded library system, run the CSADMIN utility on all other servers, in any order, where the library system has content search repositories. Invoke the utility as in step 2, responding 'y' at each CSADMIN prompt to rebuild all the content search repositories of the library system on the server.

Each set of library services on a server has its own Verity K2 Server and one or more K2 Brokers on that server to manage indexing operations. In turn, all the K2 Servers on a server are governed by one K2 Master Admin, which may be on the same server or some other server. In general, a Master Admin may govern K2 Servers for one or more library systems, whose initial and additional services are spread over several servers.

Rebuilding a content search repository on a server should attach it to a K2 Server on the same server. To verify that the rebuilt content search repositories are attached, run `csk2cfg <dbserver> <libname>` on each server.

Upgrading Content Services Unattended

As an alternative to running the interactive Content Services Setup program, you can run a non-interactive (that is, unattended) Content Services installation program to upgrade initial and additional library system services.

Just as for an interactive upgrade, you need to do some preparation before beginning an unattended upgrade of Content Services. Perform the following tasks in the order shown:

1. Refer to “Upgrade Overview” on page 152 and “Upgrade Constraints” on page 153, for information on what an upgrade involves.
2. Refer to the “Upgrade Roadmap” on page 75 for a complete checklist of the tasks required when upgrading a library system.

To perform an unattended upgrade, you need to do the following:

1. Prepare a parameter file to provide input to the Content Services installer program (refer to “Preparing a Parameter File” on page 137).
2. Copy the complete installation media to a local physical drive.
3. Invoke the Content Services installer program using the parameter file as input to the program (refer to “The dssetup Command Syntax and Options” on page 138 for command options and syntax).
4. Rebuild the content search repositories for the upgraded library systems (refer to “Rebuilding Content Search Repositories” on page 160).

The Content Services installer program used for unattended upgrade of initial or additional library services is named **dssetup.exe** and is on the first *FileNet Content Services* CD-ROM.

Preparing a Parameter File

The parameter file providing input to **dssetup** can be generated in either of two ways:

- Make a copy of the sample parameter file for upgrading initial library services or additional library services from the first *FileNet Content Services* CD-ROM and edit it to meet the needs of your organization.
- Run the (interactive) Content Services Setup program as far as the Review Setup Entries screen. But click Cancel, instead of Next, to exit from CS Setup without installing any software. The parameter file containing your installation selections is **setup.inf**, located at *InstallDrive:\filenet*.

Note If your original installation of Content Services is earlier than CS 5.4, the **setup.inf** file is located at *InstallDrive:\idmids* or *InstallDrive:\saros*.

The sample parameter files for unattended upgrades of initial library services and additional library services are available on the first *FileNet Content Services* CD-ROM, as shown in the following table.

| Sample Parameter File | Installation Task |
|-------------------------------|-------------------------------------|
| sample_upgrade.inf | Upgrade initial library services |
| sample_upgradeaddl.inf | Upgrade additional library services |

Using a Parameter File to Upgrade Library Services

Once you have prepared a parameter file, as explained above, you are read to run **dssetup** to perform an upgrade of initial or additional library services:

To perform an unattended library system upgrade:

1. At a DOS command prompt invoke the **dssetup** command to perform the installation. For example, if the sample file is named **upgr_setup.inf**, run the following command:

```
dssetup -f C:\upgr_setup.inf -g NO
```

This command will perform an installation according to the parameters specified in **C:\upgr_setup.inf** without prompting you for inputs. Refer to [“Parameter File Format and Content” on page 139](#) for command-line syntax and options.

2. If this is the first upgrade of library services on the server, reboot the server when **dssetup** exits.

After upgrading every set of services for each library system on all the servers, continue at [“Rebuilding Content Search Repositories” on page 160](#).

Separating a Database from Library Services

If you have installed (or upgraded) initial or additional library services on a server where your Oracle database is also installed, you can move the database to a separate server using the procedure in [“Moving the Database”](#) on [page 352](#).

Configuring Content Services

After you have installed the library system software, you will need to configure the library system to meet the specific requirements of your site. We recommend that you use the checklist provided in “[Configuration Roadmap](#)” on page 72.

The topics in this section of the online help describe the settings you will likely be concerned with during the initial configuration of a library system. These settings are based on properties in the System object that provide system-wide default values or control mechanisms for various system functions. Refer also to the “[Administration Guide](#)” on page 222 for additional management and configuration information.

The CS Configuration Tool



The CS Configuration tool (available in the Windows Control Panel) allows you to view or edit many of a library system's configuration parameters.

- “[Installation Tab](#)” on page 165 allows you to view (not edit) many of the installation parameters for a selected library system.
- “[Storage Manager Tab](#)” on page 167 allows you to view or edit many of the Storage Manager configuration settings for a selected library system.
- “[Replication Tab](#)” on page 169 allows you to view or edit many of the Replication Services configuration options for a selected library system.
- “[Windows User Import Tab](#)” on page 178 provides GUI controls for a tool you can use to import users and groups from your Windows network domains into your library systems.
- “[LDAP User Import Tab](#)” on page 180 provides GUI controls for a tool you can use to import users and groups from an LDAP directory into your library systems.
- “[XML User Import Tab](#)” on page 182 provides GUI controls for a tool you can use to import users and groups from an XML file into your library systems.

If you are running Content Services in a cluster environment, be aware that any modifications you make to library configuration parameters using CS Configuration will appear only on the active node of the cluster. To maintain parameter consistency among the nodes in the cluster, use Cluster Administrator to failover to the passive node. Then make the same modifications to the library parameters on that node.

Installation Tab



If your server has components for multiple library systems installed, you can select the library system whose configuration you want to view. All other fields on this tab are read-only; they display configuration parameters that were set during installation and cannot be changed.

| Field | Description |
|----------------------------|---|
| Database type | The type of database management system you are using for all library systems on this server. |
| CS version | The version number of the Content Services software currently installed on this server. |
| Root directory | The top-level directory where the Content Services software is installed. |
| Active Directory Active | The value is “Yes” if you selected the “Enable Active Directory lookups” check box in the Content Services Manager applet while logged on as a domain user. |

Library Systems

This section of the tab displays installation parameters specific to the selected library system.

| Field | Description |
|-----------------|--|
| Library system | Select the name of the library system whose configuration you want to view. |
| Storage manager | The host name and drive letter on which the initial Storage Manager is installed. |
| Content search | The host name and drive letter on which the initial Content Search Manager is installed. |
| Replication | The initial storage server host name and root directory where the Replication daemon for this library system is installed. |

| Field | Description |
|------------------------|---|
| Directory Server Type | <p>These options specify how users and groups to be imported into a library system are determined. Refer to “Importing Users and Groups” on page 272 for details.</p> <ul style="list-style-type: none"> • WinNT/Win2000 If selected, users and groups are determined by membership in a Windows NT domain or Windows 2000 or Windows 2003 Active Directory. • LDAP 3.0 Server If selected, users and groups are determined by membership in an LDAP directory on a server that supports version 3.0 (or higher) of LDAP. • FileNet XML File If selected, users and groups are determined by membership in an LDAP directory and placed in an XML file by the LDAP_EXP utility (see “LDAP_EXP” on page 435). |
| Content Search | <p>Engine char set: The character set used during indexing and searching.</p> <p>Engine char filt: The filter used during indexing.</p> |
| Content Index Creation | <p>Engine normalize: The normalization setting used during content search repository creation.</p> <p>Engine stop file: The Engine stop file used during content search repository creation.</p> |
| Content Search Filter | <p>Filter char set: The character set used in filtering.</p> <p>Filter file format: The file format expected during filtering.</p> |

Storage Manager Tab



This tab allows you to view or edit many of the Storage Manager configuration options for a selected library system.

Note If you make changes to the data in any of the editable fields, you must click **Apply** and then stop and restart the Storage Managers for your changes to take effect. Refer to “[Controlling Library System Services](#)” on page 198 for information about stopping, starting, and checking the status of services.

| Field | Description |
|--------------------|---|
| Library system | The name of the library system whose configuration you want to view or edit. (You can select the library only in the Installation tab of the CS Configuration tool.) |
| Property server | The host name of the server where the database that supports the library system is installed. This value is set when you install the library system and cannot be edited with the CS Configuration tool. |
| Log file | The path and file name of the circular log of Storage Manager operations. This file will grow no larger than the specified Max file size before looping back and overwriting itself. By default, no log file is specified. When the Storage Manager is stopped and restarted, the existing log file is renamed by appending a .prev suffix, and a new, empty log file is created with the default log file name. At most one .prev log file can exist. If you stop and restart the Storage Manager again, the Storage Manager will rename the existing log file as the .prev log file and create a new, empty log file with the default log file name. |
| Max file size | Indicates the maximum number of kilobytes of processing information the circular log file will hold before looping back and overwriting itself. Default is 256. |
| Log SQL statements | If this box is checked, all generated SQL calls will be logged. Default is not to log SQL calls. |
| Log console | If this box is checked, logging information is written directly to the console window associated with the process. Default is not to send logging information to the console. |
| DB heart beat | Seconds: The interval for checking whether the library system database is running. Default is 300. Offset: A modifier for offsetting the start time of DB heart beat interval. Default is 0 seconds. |
| DB threshold | Seconds: The interval for checking whether the database has reached its maximum disk-space threshold. Default is 300. Offset: A modifier for offsetting the start time of DB threshold. Default is 38 seconds. |

| Field | Description |
|----------------------|---|
| DB error log | <p>Seconds: The interval for checking for errors posted in the database error log file. Default is 300.</p> <p>Offset: A modifier for offsetting the start time of DB error log. Default is 75 seconds.</p> |
| Reclaim check | <p>Seconds: The interval at which the reclaim daemon runs to process queued reclaim requests. Default is 300.</p> <p>Offset: A modifier for offsetting the start time of Reclaim Check. Default is 112 seconds.</p> |
| CS index | <p>Seconds: The interval at which the Content Search daemon runs. Default is 300.</p> <p>Offset: A modifier for offsetting the start time of CS Index. Default is 150 seconds.</p> |
| Client threads | The number of client threads/processes to start. Default is 20. |
| Auto archive delay | Delays autoarchiving of versions following a checkin. May improve performance of some user interfaces by eliminating database update contentions. The integer value you enter is the number of seconds of delay. Default is 5. |
| Max clients waiting | The maximum number of clients allowed to wait in the worker queue. Default is 8. |
| Max connects waiting | The maximum number of clients allowed to wait in the communications listening queue. Default is 5. |
| Secure Deletes | <p>These options control the Secure Document Delete feature. Refer to “Secure Document Delete” on page 331 for details.</p> <ul style="list-style-type: none"> • No secure deletes <p>Ordinary delete operation is performed; files are not overwritten before being deleted.</p> • One scrub <p>Files are overwritten with zeroes once before being deleted. This level corresponds to the <i>Clear</i> definition of the DOD 5220.22-M specification.</p> • Three scrubs <p>Files are overwritten once with an arbitrary character, again by the character's complement, and finally by a random character before being deleted. This level corresponds to the <i>Purge</i> or <i>Sanitize</i> definition of the DOD 5220.22-M specification.</p> |

Replication Tab



You use the CS Configuration Tool Replication tab to control replication configuration options for a library system.

Note If you make changes to the data in any of the editable fields, you must click **Apply** and then stop and restart Replication Services for your changes to take effect. Refer to “[Controlling Library System Services](#)” on page 198 for information about stopping, starting, and checking the status of services.

| Field | Description |
|-----------------|--|
| Error alert to | The computer name where an alert message is sent (using NetMessageSend) when an error occurs. The default is an empty field. |
| Wakeup seconds | The approximate latency period (in seconds) the Replication Daemon uses before beginning to process replication requests for the associated Replication Services. The default is 30 seconds. Note You cannot set this value to 0, so that Replication Daemon runs continuously. |
| Log file | The name of the replication log file stored in the <code>\filenet\repl\logs\librarysystem</code> subdirectory. The default file name is repl.log. If you do not enter a valid file name, the log file is not created. |
| Max size (K) | The maximum size (in kilobytes) of the replication log file. The default is 128. |
| Truncate to (K) | The size (in kilobytes) used for truncating the replication log file when the On Startup or When Full options are set to TRUNCATE. The default is 16. |
| Error threshold | The number of times a replication request will be attempted before logging a failure and ceasing attempts to resolve the request. The default is 5. |
| On startup | Determines the action the Replication Daemon performs when opening the replication log. Select one of the following values (default is APPEND): <ul style="list-style-type: none"> • APPEND Adds the new log entries to the end of the current replication log file. • EMPTY Deletes the previous entries in the log. • TRUNCATE Removes the oldest entries in the log until the replication log is no greater than the size specified in the Truncate To option. |

| Field | Description |
|---------------|---|
| When full | <p>Determines the action taken when the replication log file is full, based on the Max Size setting.</p> <p>Select one of the following values (default is TRUNCATE):</p> <ul style="list-style-type: none"> • EMPTY Deletes the previous entries and resumes logging. • STOP Ends replication logging. • TRUNCATE Removes the oldest entries in the log until the replication log is no greater than the size specified in the Truncate To option and resumes logging. |
| Logons | <p>Determines whether login attempts to the library systems on this property server are automatically logged in the replication log file.</p> <p>Select one of the following values (default is 5):</p> <ul style="list-style-type: none"> • ALWAYS All login attempts are logged. • ERROR Only login attempts resulting in a failure are logged. • N All login attempts after the specified number (<i>N</i>) of consecutive failed login attempts are logged. The default is 5. • NEVER No login attempts are logged. |
| Full alert to | <p>The computer name where an alert message is sent (using NetMessageSend) when the replication log file becomes full, based on the MaxSize parameter.</p> <p>The default is an empty list.</p> |
| Requests | <p>Determines whether the creation of work requests for the target libraries is automatically logged. Work requests are created on the original library.</p> <p>Select one of the following values (default is ERROR):</p> <ul style="list-style-type: none"> • ALWAYS All work requests are logged. • ERROR Only failed work requests are logged. • NEVER No work requests are logged. |

| Field | Description |
|-----------------|--|
| Targets | <p>Determines whether the processing of work requests by the target libraries is automatically logged. Work requests are executed by the target library.</p> <p>Select one of the following values (default is ERROR):</p> <ul style="list-style-type: none"> • ALWAYS All work requests are logged. • ERROR Only failed work requests are logged. • NEVER No work requests are logged. |
| API | <p>Determines whether library system SQL log messages are written to the replication log file.</p> <p>Select one of the following values (default is NEVER):</p> <ul style="list-style-type: none"> • ALWAYS All log messages are logged. Use this setting only when troubleshooting because it can expand the replication log file considerably. • NEVER No log messages are logged. |
| Library | <p>The name of the library system to be managed by the Replication Daemon. Select the library from the library-system drop-down list in the Installation tab of the CS Configuration tool.</p> |
| Remote interval | <p>Specifies how often (in minutes) the Replication Daemon checks and updates replicas on this library system. The default is 5 minutes.</p> <p>Note You cannot set this value to 0, so that Replication Daemon runs continuously.</p> <p>This setting is used in conjunction with the time ranges specified using the Run Times button.</p> |

| Field | Description |
|----------------------|--|
| Index replicas | <p>Determines whether document replicas are indexed for content search on the target library system.</p> <p>Select one of the following values (default is ITEM):</p> <ul style="list-style-type: none"> • DEFAULT Use the Default Index Versions setting from the target library system to determine whether to index replicas. • ITEM Determines whether to index the replica based on the Index setting from the version being replicated. • NO Never index replicas on this library system. • YES ALWAYS index replicas on this library system. |
| Replica online limit | <p>Determines the online limits of document replicas on the target library system.</p> <p>Because replicas are exact copies of the original document, we recommend using the default setting of 1.</p> <p>Select one of the following values (default is 1):</p> <ul style="list-style-type: none"> • DEFAULT Use the Online Limit setting for this library system as the setting for replicas. • ITEM Use the Online Limit setting from the original document as the setting for replicas. • <i>N</i> Use this exact number as the Online Limit setting for document replicas. |

| Field | Description |
|---------------------|--|
| Replica index limit | <p>Determines the index limits used for document replicas on this library system.</p> <p>Because replicas are exact copies of the original document, we recommend using the default setting of 1.</p> <p>Select one of the following values (default is 1):</p> <ul style="list-style-type: none">• DEFAULT Use the Index Version Limits setting for this library system as the setting for replicas.• ITEM Use the Index Version Limits setting from the original document as the setting for replicas.• <i>N</i> Use this exact number as the Index Version Limit setting for document replicas. |

| Field | Description |
|-------------------|---|
| Item title format | <p>Determines the display text for the document title on the target library system. To help your end users identify the original library for a document, you can include the Library constant (%l) in the document title. Or refer to the special icon IDM Desktop uses to indicate replica documents.</p> <p>The default (%t) displays the document title. You can specify one or more of the following constants in any combination:</p> <ul style="list-style-type: none"> • %l Inserts the name of the original library system for the original document. • %o Inserts the document ID of the original document. • %s Inserts the name of the property server for the library system in which the original document is located. • %t Inserts the title of the original document. Replication Services always appends this constant if it is not already specified in the format string. • %% Inserts a percent (%) character. This constant allows you to include a literal percent character in your format string. Replication Services treats any other characters not preceded by a percent character as literals. <p>Any additional characters you specify will also appear in the title.</p> <p>If quotes (single or double) appear at the beginning and end of the format string, a space is added to the beginning and end of the title.</p> <p>Examples:</p> <p><code>[%l]%t</code> Displays as: [Library]Title</p> <p><code>[%s:%l]%t</code> Displays as: [Server:Library]Title</p> <p><code>%t-> Replica</code> Displays as: Title-> Replica</p> <p><code>Alt+0174%t</code> (press the Alt key and type 0174 on the numeric keypad) Displays as: ®Title</p> |

| Field | Description |
|---------------------|--|
| Folder title format | <p>Determines the display text for the folder title on the target library system. To help your end users identify the original library for a folder, you can include the Library constant (%l) in the folder title.</p> <p>The default (%t) displays the folder title. You can specify one or more of the following constants in any combination:</p> <ul style="list-style-type: none"> • %l Inserts the name of the original library system for the original folder. • %s Inserts the name of the property server of the library system in which the original folder is located. • %t Inserts the title of the original folder. Replication Services will append this constant if it is not already specified in the format string. This action ensures that Replication Services cannot inadvertently create replicas having improper identification. • %% Inserts a percent (%) character. This allows you to include a literal percent character in your format string. Replication Services treats any other characters not preceded by a percent character as literals. <p>Any additional characters you specify will also appear in the title.</p> <p>If quotes (single or double) appear at the beginning and end of the format string, a space is added to the beginning and end of the title.</p> <p>Examples:</p> <p><code>[%l]%t</code> Displays as: [Library]Title</p> <p><code>[%s:%l]%t</code> Displays as: [Server:Library]Title</p> <p><code>%t-> Replica</code> Displays as: Title-> Replica</p> <p><code>Alt+0174%t</code> (press the Alt key and type 0174 on the numeric keypad) Displays as: ®Title</p> |
| Enable Replication | When checked, Replication Services is enabled. |

| Field | Description |
|---------------------|---|
| Replicate item ACLs | <p>Determines whether the access list (ACL) for a document replica is based on the access list of the original document. The default is Checked.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> • Checked Inside an administrative domain, the access list for a document replica contains the ReplDaemon user with Owner access rights and the access list entries from the original document (with no entry having greater than Author access rights). Outside an administrative domain, a document replica receives an access list containing the ReplDaemon user with Owner access rights and the access list entries from the original document (with no entry having greater than Viewer access rights). However, the users and groups from the ACL of the original document must also exist on the target library system. If a user or group does not exist on the target library system, the user or group entry is not included in the document replica access list. • Unchecked Inside an administrative domain, a document replica receives an access list containing the target library system's ReplDaemon user with Owner access rights and the ReplDaemon user's Default Item Access List entries. Using the Content Services Admin Tools, you can modify the default item access list in the ReplDaemon's User object profile. No document replica ACL will have greater than Author access rights. Outside an administrative domain, a document replica receives an access list containing the target library system's ReplDaemon user with Owner access rights and the ReplDaemon user's Default Item Access List entries. Using the Content Services Admin Tools, you can modify the default item access list in the ReplDaemon's User object profile. No document replica ACL will have greater than Viewer access rights. |

| Field | Description |
|----------------------------|--|
| Replicate folder ACLs | <p>Determines whether the access list (ACL) for a folder replica is based on the access list of the original folder. The default is Checked.</p> <p>Select one of the following values:</p> <ul style="list-style-type: none"> Checked Inside or outside an administrative domain, a folder replica receives an access list containing the ReplDaemon user with Owner access rights and the access list entries from the original folder (with no entry having greater than Author access rights). Outside an administrative domain, the users and groups from the ACL of the original folder must also exist on the target library system. If a user or group does not exist on the target library system, the user or group entry is not included in the folder replica access list. Unchecked Inside or outside an administrative domain, a folder replica receives an access list containing the ReplDaemon user with Owner access rights and General Users with Viewer access rights. |
| Replicate storage category | <p>Determines whether the storage category assigned to document replicas on the target library system is the same as the original document's storage category. The storage category must already exist on the target library system. It will not be replicated from the original library system.</p> <p>The default is Checked.</p> <p>If you uncheck this option, Replication Services uses the ReplDaemon user's default storage category for the replicas. Using the Content Services Admin Tools, you can modify the default storage category in the ReplDaemon's User object profile.</p> |
| Replicate archive category | <p>Determines whether the archive category assigned to document replicas on the target library system is the same as the original document's archive category. The archive category must already exist on the target library system. It will not be replicated from the original library system.</p> <p>The default is Checked.</p> <p>If you uncheck this option, Replication Services uses the ReplDaemon user's default archive category. Using the Content Services Admin Tools, you can modify the default archive category in the ReplDaemon's User object profile.</p> |
| Run Times | <p>Click the Run Times button to specify the time range during which the Replication Daemon synchronizes replicas on this library system with the original objects.</p> <p>The default is the range 00:00-23:59 for each day of the week.</p> |

Windows User Import Tab



This tab provides GUI controls by which you can enable one-way synchronization of your library system groups and users with your Windows network domain groups and users. This tab is available on the initial storage server only (not on additional storage servers).

Note If you make changes to the data in any of the editable fields, you must click **Apply** and then stop and restart the Storage Manager for your changes to take effect. Refer to [“Controlling Library System Services” on page 198](#) for information about stopping, starting, and checking the status of services.

| Field | Description |
|--------------------------|--|
| Library system | The library system whose user-import configuration you want to view or edit. (You can select the library only in the Installation tab of the CS Configuration tool.) |
| Enable user import | Select this check box to enable the import daemon to run. Clear the check box to shut down the daemon. Default is not to run the daemon. |
| Create preview file only | Select this check box to enable generation of a preview file showing the transactions that would be performed during a user import. Clearing the check box results in importing the groups/users instead of generating a preview file. |
| Interval (seconds) | Sets the interval (in seconds) between successive executions of the import daemon. The default is hourly (3600 seconds). |
| Offset (seconds) | An offset (in seconds) for the starting time of the import daemon with respect to other processes that run under the Storage Manager. The default is 25 minutes (1500 seconds). See “Storage Manager Tab” on page 167 for information about the other Storage Manager processes with configurable offsets. |
| Log file | The name of the file for the log produced when groups/users are imported. The default path and name of the log file is Drive:\filenet\LibrarySystem\dirsync.log |
| Log size | The maximum size (in kilobytes) to which the log file can grow before overwriting itself. This field is blank by default. |
| Preview file | The name of the transaction file that shows the individual operations that would occur if the users/groups are actually imported. The default name of the preview file is preview.log . |
| Group List File | The name of the file containing the list of groups being imported to Content Services. The default path of this file is InstallDrive:\filenet\LibrarySystem\group.log . |
| Domain | The NT domain containing the groups and users you want to import into your library system. |
| Available Groups | The groups in the NT domain that you can select for importing into a library system. |

| Field | Description |
|-----------------|---|
| Selected Groups | The names of the groups selected for importing into a library system. |

LDAP User Import Tab



This tab provides GUI controls by which you can enable one-way synchronization of your library system groups and users with the groups and users stored in an LDAP Directory Server. This tab is available on the initial storage server only (not on additional storage servers).

Note If you make changes to the data in any of the editable fields, you must click **Apply** and then stop and restart the Storage Manager for your changes to take effect. Refer to [“Controlling Library System Services”](#) on page 198 for information about stopping, starting, and checking the status of services.

| Field | Description |
|--------------------------|---|
| Library system | The library system whose user-import configuration you want to view or edit. (You can select the library only in the Installation tab of the CS Configuration tool.) |
| Enable user import | Select this check box to enable the import daemon to run. Clear the check box to shut down the daemon. Default is not to run the daemon. |
| Create preview file only | Select this check box to enable generation of a preview file showing the transactions that would be performed during a user import. Clearing the check box results in importing the groups/users instead of generating a preview file. |
| Interval (seconds) | Sets the interval (in seconds) between successive executions of the import daemon. The default (for NDS directories, but not iPlanet directories) is hourly (3600 seconds). |
| Offset (seconds) | An offset (in seconds) for the starting time of the import daemon with respect to other processes that run under the Storage Manager. The default (for NDS directories, but not iPlanet directories) is 25 minutes (1500 seconds). See “Storage Manager Tab” on page 167 for information about the other Storage Manager processes with configurable offsets. |
| Log file | The name of the file for the log produced when groups/users are imported. The default path and name of the log file (for NDS directories, but not iPlanet directories) is Drive:\filenet\LibrarySystem\dirsync.log |
| Log max size (kb) | The maximum size (in kilobytes) to which the log file can grow before overwriting itself. The default (for NDS directories, but not iPlanet directories) is 256 KB. |
| Preview | The name of the transaction file that shows the individual operations that would occur if the users/groups are actually imported. The default name of the preview file is preview.log . |
| Group List File | The name of the file containing the list of groups being imported to Content Services. The default file name is group.log . |
| LDAP Server:Port | The LDAP server name and port number (if the port is not the default port) in the form server:port. |

| Field | Description |
|---------------------------------|---|
| BaseDN | The initial portion of a distinguished name (user) logging on to the LDAP directory. |
| BindDN | The distinguished (user) name used to log on to the LDAP directory. |
| BindDN Password | The password for the bindDN user |
| Available Groups / OUs | The organizational units (groups) in the LDAP directory you can select for importing into a library system. |
| Selected Groups / OUs | The names of the organizational units (groups) selected for importing into a library system. |
| Refresh | Clicking this button, after specifying the LDAP Server:Port, BaseDN, BindDN, and BindDN Password, triggers a login to the LDAP directory and will cause all the available groups or organizational units to appear in the Available Groups / OU box if you have specified values for LDAP Server:Port, BaseDn, BindDN, etc. |
| Select for Organizational Units | If this checkbox is cleared, groupOfNames and groupOfUniqueNames object classes are used to find groups to import. If this checkbox is selected, then LDAP organizational units are used. |

XML User Import Tab



This tab provides GUI controls by which you can enable one-way synchronization of your library system groups and users with the groups and users specified in an XML file. This tab is available on the initial storage server only (not on additional storage servers).

Note If you make changes to the data in any of the editable fields, you must click **Apply** and then stop and restart the Storage Manager for your changes to take effect. Refer to “[Controlling Library System Services](#)” on page 198 for information about stopping, starting, and checking the status of services.

| Field | Description |
|--------------------------|---|
| Library system | The library system whose user-import configuration you want to view or edit. (You can select the library only in the Installation tab of the CS Configuration tool.) |
| Enable user import | Select this check box to enable the import daemon dirsync.exe to run. Clear the check box to shut down the daemon. Default is not to run the daemon. |
| Create preview file only | Select this check box to enable generation of a preview file showing the transactions that would be performed during a user import. Clearing the check box results in importing the groups/users instead of generating a preview file. |
| Interval (seconds) | Sets the interval (in seconds) between successive executions of the import daemon. The default (for NDS directories, but not iPlanet directories) is hourly (3600 seconds). |
| Offset (seconds) | An offset (in seconds) for the starting time of the import daemon with respect to other processes that run under the Storage Manager. The default (for NDS directories, but not iPlanet directories) is 25 minutes (1500 seconds). See “ Storage Manager Tab ” on page 167 for information about the other Storage Manager processes with configurable offsets. |
| Log file | The name of the file for the log produced when groups/users are imported. The default path and name of the log file (for NDS directories, but not iPlanet directories) is Drive:\file\LibrarySystem\dirsync.log |
| Preview log file | The name of the transaction file that shows the individual operations that would occur if the users/groups are actually imported. The default name of the preview file is preview.log . |
| Log max size (kb) | The maximum size (in kilobytes) to which the log file can grow before overwriting itself. The default (for NDS directories, but not iPlanet directories) is 256 KB. |
| Group list file | The name of the file containing the list of groups being imported to Content Services. The default file name is group.log . |
| XML file | The fully qualified path to the XML file that contains the users and groups to be imported into a CS library system. The file must <i>not</i> reside on a mapped drive. |

| Field | Description |
|----------------------|---|
| Available XML Groups | The groups in the LDAP directory you can select for importing into a library system. |
| Selected XML Groups | The names of the groups selected for importing into a library system. |
| Refresh | Clicking this button, after specifying an XML file, will trigger a logon to the LDAP directory and will cause all the available groups to appear in the Available XML Groups box. |

Imported Group Defaults

The Group Defaults dialog box allows you to set property values for new groups you are adding to your library system via the User Import tool.

Note Changing the values in this dialog box will not update the property values of existing library system groups.

To display this dialog box, highlight a group name in the Selected Group list on the User Import tab, then right mouse click and select Group Defaults. These settings allow you to control the default values for the Group object properties listed below. Refer to the “Objects and Properties Reference” in the *CS Explorer Online Help* for more information.

| Field | Default Value | Definition |
|--------------------------|---------------|---|
| Default storage category | Documents | <p>Provides an initial value for the Storage Category property when new Item objects are added. When new items are added, the system checks the following objects in the order shown for this property value and uses the first value encountered as the storage category for the new Item:</p> <ol style="list-style-type: none"> 1. User object 2. User's active Group object 3. System object |
| Default archive category | Archive | <p>Provides an initial value for the Archive Category property when new Item objects are added. When new items are added, the system checks the following objects in the order shown for this property value and uses the first value encountered as the archive category for the new Item:</p> <ol style="list-style-type: none"> 1. User object 2. User's active Group object 3. System object |

| Field | Default Value | Definition |
|-------------------------|---------------|--|
| Primary storage manager | Blank | <p>Indicates the Primary Storage Manager for the system to use first when handling an operational request generated by a member of this group (for example, check out or check in). The system checks the following objects in the order shown for this property value:</p> <ol style="list-style-type: none"> 1. User object 2. User's active Group object 3. System object <p>The library system attempts to use the first Primary Storage Manager value it finds. If this Storage Manager is unavailable, the system continues to check all the Storage Managers (in alphabetical order) until it finds one to handle the request.</p> |
| Default path | Blank | <p>Sets the default path for delivery of checked out or copied version items.</p> |
| Access domain | Blank | <p>Controls the setting of the group's Access Domain property, which indicates the group affiliations to apply when determining the search results to return. Valid values are:</p> <ul style="list-style-type: none"> • Active Group: Sets search results access rights to the user's active group. • All Groups: Sets the search results on the highest access rights of all groups in which the user is a member. • Blank: Refers search results to the System object. <p>The library system checks the Active Group object first and then the System object and accepts the first value it finds. However, the value is overridden by any value specified in the User's properties.</p> |

| Field | Default Value | Definition |
|--------------------------|---------------|---|
| Default display language | Blank | <p>Select English (default), French, or German from the drop-down menu.</p> <p>Before attempting to configure a non-English display language, please refer to the technical notice titled, <i>Installing Content Services into a Non-English Environment</i> available on the FileNet Worldwide Customer Support web site. See Documentation Updates under “About the Documentation” on page 12 if you need help.</p> |
| Content search threshold | Blank | <p>Determines whether to do a content search or property search first, when both types of criteria are specified for a search. This threshold indicates that if the number of content search hits returned by a search is greater than the threshold value, the property-based search criteria are used first. If this value is less than or equal to the number of content search hits, the content search-based criteria are used first. If a Content Search Threshold property is not found, the default threshold value used is 100.</p> |
| Secure Searches | Blank | <p>This property indicates whether an active group member must have at least Viewer access rights to an item to have that item included in the results of a search. Select Yes, No, or Blank (the default) from the drop-down menu.</p> <ul style="list-style-type: none"> • The <i>No</i> value indicates that all items meeting the search criteria will be returned for a search, regardless of the user's access rights to those items. (However, the user must still have at least Viewer access rights to an item to see its associated Item object.) • The <i>Yes</i> value indicates that only those items to which a user has at least Viewer access rights will be returned in the results of a search. • A <i>Blank</i> value defaults to the Secure Searches? property value in the System object. |
| Default index versions | Checked | <p>Specifies a value for the Default Index Versions? property, which indicates whether versions of an item are indexed by default.</p> <p>If this box is checked, this property is set to <i>Yes</i> for the imported Group object. If this box is cleared, the value of this property in the System object is used.</p> |

| Field | Default Value | Definition |
|------------------|---------------|--|
| Add items | Checked | When checked, permits imported users to add items to the library system. |
| Add users | Blank | When checked, permits imported users to add users to the library system. |
| Add groups | Blank | When checked, permits imported users to add groups to the library system. |
| No Group Updates | Blank | <p>If this box is checked, a group imported into a library system will not have its property values changed by a subsequent import of this group, even if the default values for group imports have changed since the original import.</p> <p>When the group is imported, the Imported Group? property will automatically be set to NLTM_NOUPDATE, XML_NOUPDATE, or LDAP_NOUPDATE, depending upon which directory service was used for synchronization.</p> |
| No User Updates | Blank | <p>If this box is checked, a user imported into a library system will not have its property values changed by a subsequent import of this user, even if the default values for the group to which the user belongs have changed since the original import.</p> <p>When the user is imported, the Imported User? property will automatically be set to NLTM_NOUPDATE, XML_NOUPDATE, or LDAP_NOUPDATE, depending upon what directory service was used for synchronization.</p> |
| Group enabled | Checked | If this box is checked, the imported users will be able to change to this group or open a session using this group as their active group; if cleared, they cannot. |
| Users enabled | Checked | If this box is checked, the imported user will be able to log on to the library system; if cleared, the user will be imported as disabled, and thus unable to log on to the library system. |

Understanding Default Values and the Inheritance Path

Numerous properties in the library system (especially those in the Item object) have their initial values set according to an inheritance path from properties in other objects, including the System object. This inheritance path means, for example, that when a new protected item is added, the library system takes the following steps to determine some of the initial values for the new Item object:

1. Checks the User object of the person adding the item and applies any default values it finds to the Item object.
2. Checks the user's active Group object for any default values that were not found in the User object.
3. Checks the System object for any default values that were not found in either the User or Group object.
4. Uses the library system program defaults for any default values that were not found in the User, Group, or System object.

For example, if a user named Bob, whose active group is Marketing, adds a new Item, the storage category for the new Item object will receive an initial value based on a check of the Default Storage Category properties in the following objects: (1) Bob's User object, (2) Marketing's Group object, and (3) the System object. The first of these properties to contain a value will provide the necessary default for the Item object. (In this case, the System object must contain a non-blank value, so the final check for a library system program default is never required.)

The following table shows the properties in the Item object that are affected by this inheritance path method for applying defaults:

| Inheritance Path Property | Item Property that Inherits |
|----------------------------|-----------------------------|
| Default Archive Category | Archive Category |
| Default Storage Category | Storage Category |
| Default Checkout File Name | Default Checkout File Name |
| Default Item Access List | Access List |
| Default Index Versions? | Default Index Versions? |

The following table shows properties that are not in the Item object but are affected by the same inheritance path method for applying defaults from the User, Group, and System objects:

| Inheritance Path Property | Description |
|---------------------------|--|
| Primary Storage Manager | Applies when the library system must perform an operation requiring a Storage Manager. The first value found during this check determines the storage server whose Storage Manager the library system will attempt to use to process a user's operational request (for example, checking out a version). |

| Inheritance Path Property | Description |
|---------------------------|---|
| Access Domain | <p>Applies when the library system must determine what portion of the total results of a search to display to a user. The first value found during this check determines whether the library system applies:</p> <ul style="list-style-type: none"> • The user's active group access rights. • The highest access rights of any of the groups the user belongs to. <p>Note This value cannot be left blank in the System object. Also, in some user interfaces, the user can set this value on a search-by-search basis.</p> |
| Secure Searches? | <p>Indicates whether a user must have at least Viewer access rights to an item (or object) to have that item (or object) included in the results of one of his or her searches.</p> <p>Note In some user interfaces, the user can set this value on a search-by-search basis.</p> |
| Display Language | <p>Determines in what language to display the library system's property names, field names, messages, and Yes/No values. If no value is specified, English is used by default.</p> |

Users at some sites may find it useful to set these defaults in their User objects (for example, so that they can control what values will be used to describe the items they add). At other sites, group leaders can take the responsibility of setting these values in Group objects so that users do not have to be concerned with such matters. Likewise, many sites may find it best to set the defaults at the System object level so that all users have the same initial values placed in their Item objects or are using the same display language. Finally, many sites may determine the best level at which to set the defaults on a property-by-property basis.

Note Users will be able to change their User object defaults only if they have access to their User object (or the particular properties that control these defaults) through the chosen user interface. Because this kind of functionality is not supported by many user interfaces, you may prefer to leave the User object property blank and set the values in the Group object instead, since many user interfaces do provide a change-group feature.

Additional default values that you can set in the System object will affect Item objects more directly (there are no comparable properties in the User and Group objects). These defaults are:

| System Property | Item Property that Directly Inherits |
|-------------------------------|--------------------------------------|
| Default Online Limit | Online Limit |
| Default Indexed Version Limit | Indexed Version Limit |

In all cases, users with Owner or Admin access rights to the item can change these values in the Item object at any time.

Setting Default Properties

The following are a few of the default properties that appear in the System object. For a more complete list of properties in the System object, refer to the *CS Explorer Online Help*.

Default Index Versions?

When a user adds a new protected item to the library system, this property value provides an initial setting for the Default Index Versions? property in the Item object. If this property value is set to Yes, the current version of the item is automatically indexed for content searching when the item is added to CS.

Default Item Access List

Like other objects in CS, the Item object has an Access List property. The value of this property is a list of users and groups with associated access rights to that particular object. The Default Item Access List provides an initial set of access list entries for each new item that is added to the library system. The access list in the Item object can then be modified as necessary.

If the library system searches the Default Item Access List properties in the User, Group, and System object and finds that they are all empty, it inserts the library system program default: the Added By User is granted Owner access rights. If this entry is included, the Item object's access list will receive the access rights assigned in the default item access list rather than those specified in the library system program default.

For example, in the System object's default item access list, you could grant Viewer access rights to Anne Blanchard. Then, Anne would be granted Viewer access rights to all new Item objects that received values from the System object's default item access list, even if she added the item herself.

Note To deny general users any access, simply exclude the General Users group from the default item access list.

Any user with Owner or Admin access rights to an Item object can override the initial values in its access list. All administrators have Admin access to all items, regardless of what access rights they are assigned in a particular access list.

Sequential Item IDs?

The library system assigns an identifying number to each item that users add to the library system. The Item ID property in the Item object enables either sequential Item IDs (if the property is set to Yes) or non-sequential (date-based) Item IDs (if the property is set to No). If not specified, the property defaults to Yes (use sequential Item IDs).

Sequential Item IDs

When using sequential Item IDs, an algorithm is applied to avoid conflicts with existing, non-sequential Item IDs. Because non-sequential Item IDs use digits 3-5 of the 9 digit Item ID (numbering left to right, as in 123456789) to represent the day of the year, the numbers 000 through 366 in positions 3-5 are not used when generating a sequential Item ID.

For example, the first sequential Item ID is 003670000. This number is incremented sequentially until the Item ID equals 009999999. The next Item ID after 009999999 is 013670000, which is incremented sequentially until 019999999. The next Item ID is 023670000 and so on until the maximum Item ID of 999999999 is reached. Thus, a total of 633,000,000 items may be assigned sequential Item IDs.

Non-Sequential (Date-Based) Item IDs

Non-sequential Item IDs are based on Julian dates, where digits 1-2 represent the decade, 3-5 the day of the year, and 6-9 the sequence number.

A typical Julian date item ID might appear as:

972860079

where 9 is the decade, 7 is the year, 286 is the day of the year (October 13), and 0079 is the sequence number. This example indicates the item was the seventy-ninth item added on October 13, 1997.

Because there are only four digits reserved for the sequence number, date-based Item IDs can be applied only to the first 9,999 items added to a particular library system on a particular day. Sequential numbering is automatically enabled beginning with the 10,000th item added.

To avoid conflicts with existing sequential Item IDs (generated by versions 4.1.3 or earlier library systems), if digits 1-2 of the next sequential Item ID are greater than or equal to the current decade, sequential Item IDs will be used. (As in previous releases of Content Services, if the number of items generated in a single day is greater than 9999, the system will automatically switch from non-sequential to sequential Item IDs.)

Note You will not be able to use non-sequential Item IDs in the year 2000 (the first two digits of the Item ID will always be greater than or equal to 00). In the year 2001, you will be able to use non-sequential Item IDs provided the next sequential Item ID is less than 010000000. Once sequential Item IDs reach 009999999, all items added to the library system in the year 2001 will have sequential Item IDs. In the year 2002, you will be able to use non-sequential Item IDs as long as the next sequential Item ID is less than 020000000.

Identifying the Primary Storage Manager

The Primary Storage Manager property indicates the storage server whose Storage Manager the library system will attempt to use first to handle an operational request. For example, when you check in a version, the library system will attempt to use the primary Storage Manager to handle the logistics of getting the file from where it is stored on the network to the storage repository associated with the version's storage category.

There is a Primary Storage Manager property in all User objects, Group objects, and in the System object. Refer to [“Understanding Default Values and the Inheritance Path” on page 188](#) for details on how the values in these objects are applied.

If the Storage Manager on the storage server indicated in this property is not available (or fails at some later time), the library system will use the Storage Manager on the next storage server, in alphabetical order from the list of storage servers identified in the library system. The library system will repeat this step as often as necessary to locate an available Storage Manager. Note that this is true only if a storage category is not assigned to a storage repository on the inaccessible server. For example, if a user is attempting to access a version associated with a storage category named Documents, and the Documents storage category is assigned to a storage repository on a downed server, the user will be notified that the operation cannot be completed.

You must have Admin access rights to the object (or be an administrator) to change the value of this property in the User and Group objects.

Setting Version Online Limits and Counts

To help ensure that the storage space for version files remains adequate, you will likely want to control how many versions can remain online for each protected item. Content Services provides an automatic archiving feature that

will move older versions offline as newer versions are created. To take advantage of this feature, you must make two decisions:

- What system-level limit for the maximum number of online versions is reasonable for the majority of the items created at your site?
- Should versions marked as Permanent be counted whenever this limit is checked (even though such versions cannot be moved offline by the automatic archiving process)?

The following properties in the System object enable you to determine initial values for controlling how online version counts are maintained.

Default Online Limit

The Item object's Online Limit property plays an important role in the library system by controlling the number of versions that are stored online. If adding a new version increases the number of online versions past the item's Online Limit property value, the "oldest" version that is available for archiving is selected for that purpose.

The "oldest" version is the online version with the smallest version number and Last Access Date older than all other online versions. The Last Access Date of a version is updated whenever the version is checked out, checked in, copied, or viewed.

When a version is archived, the library system sets the Availability Status property in the Version object to Offline. It then moves the document version to the archive repository associated with the Archive Category property value in the Item object. This process, referred to as automatic version archiving, helps keep the amount of disk space occupied by the versions under control.

The Default Online Limit property in the System object provides an initial value for the Online Limit property in new Item objects. For example, suppose you set the value of the Default Online Limit property to 4. Then, whenever a user adds a new protected item, the Online Limit property in the new Item object would be set to 4 by default.

Note Increasing the value of the item's online limit does not automatically restore versions that have been moved offline.

Users with Owner or Admin access rights to the Item object can override the initial value of the Online Limit property, but it serves as a guideline. To make this guideline as effective as possible, look at the informational needs of the users and try to find out how many online versions are necessary to carry out day-to-day activities. For example, there may be no need for versions other than the current version to reside online for most items. In this case, you would set the Default Online Limit property value to 1. Then, users could override it on an item-by-item basis when necessary.

Permanent Versions in Limit?

The Permanent Versions in Limit? property determines whether versions marked as permanent are to be included in the online limit for a protected item. The default value for this property is Yes.

If you set the value of this property to No and a particular item is at its online limit, this item could exceed the limit if you were to change the value to Yes. In this situation, the library system would present the standard over-the-limit message to the next user who checks in a version.

Setting Version Index Limits and Counts

(For library systems with a Content Search Managers only.) Just as there is good reason to limit the number of versions stored online, there is equally good reason to limit the number of versions indexed for content searches: it will help to conserve disk space and improve performance. If too many versions are indexed, you will end up using more space and the library system will respond less quickly when searching for information.

The following properties in the System object enable you to set an initial indexed version limit for new items and indicate whether permanently indexed versions should be counted whenever this limit is checked.

Default Indexed Version Limit

The Default Indexed Version Limit property in the System object provides an initial value for the Indexed Version Limit property in new Item objects. The latter property, in turn, plays an important role in the library system by controlling the number of versions that can be indexed for content searching. If indexing a version increases the number of indexed versions past the item's Indexed Version Limit property value, the “oldest” indexed version is selected for deindexing. (The “oldest” version is the version with the lowest version number.)

For example, suppose you set the value of the Default Indexed Version Limit property in the System object to 4. Then, whenever a user adds a new protected item, the Indexed Version Limit property in the new Item object would be set to 4 by default.

Users with Owner or Admin access rights to the Item object can override the initial value of the Indexed Version Limit property, but it serves as a guideline. To make this guideline as effective as possible, look at the informational needs of the users and try to find out how many indexed versions are necessary to carry out day-to-day activities. For example, there may be no need for versions other than the current version of an item to be available for content searching. In this case, you would set the Default Indexed Version Limit property value to 1. Then, users could override it on an item-by-item basis when necessary.

Note Currently, versions are not deindexed automatically when they are moved offline during the archiving process. Therefore, unless your users have a frequent need to search on the contents of offline versions, you should probably make sure that the Default Indexed Version Limit value is less than, or equal to, the Default Online Limit value. These coordinated settings will ensure that the content search indexes are no larger than they have to be, since versions will be deindexed before being automatically archived.

Setting an item's Indexed Version Limit to zero still keeps one item indexed.

Permanent Indexes In Limit?

The Permanent Indexes in Limit? property determines whether versions marked as permanently indexed for content searching in their Version objects are to be included in the indexed version limit for a protected item. The default value for this property is Yes.

If you set the value of this property to No and a particular item is at its indexed version limit, this item could exceed the limit if you were to change the value to Yes. In this situation, the library system would present the standard over-the-limit message to the next user who indexes a version.

Controlling Search Results

The following properties in the System object enable you to limit or broaden the results of a search by setting security, scope, and content searching.

Setting Secure Searches

The Secure Searches? property value in the System object indicates whether a user must have at least Viewer access rights to an item (or object) to have that item (or object) included in the results of one of his or her searches. A No value indicates that all items meeting the search criteria will be returned for a search, regardless of the user's access rights to those items. (However, the user must still have at least Viewer access rights to an item to open its associated Item object.) A Yes value indicates that only those items to which a user has at least Viewer access rights will be returned in the results of a search.

Like most inheritance path properties in CS, a Secure Searches? property appears in User, Group, and System objects. Refer to [“Understanding Default Values and the Inheritance Path” on page 188](#) for details on how values are applied from these objects.

You may want to set this value to Yes if your library system is being used to store particularly sensitive material to which access must be strictly controlled. On the other hand, if you don't mind users having knowledge of (and copy access to) all system-controlled material and you want the added benefit of faster search operations, set the value to No.

The Secure Searches? property controls searches to all library system objects having access lists. For example, because folders in IDM Desktop are controlled internally by Custom objects (which have access lists), if you set the Secure Searches? property to No, then all users will be able to see all folders, even if their access rights have been set to None. Those users having None access rights, however, will be unable to perform any folder operations other than viewing, as expected.

The library system provides a default value of blank in the User and Group objects and No in the System object, which you can override using the associated toggle button.

Note Some user interfaces may allow the Secure Searches? property to be set as part of the search itself. If this value is set to Yes in the Search object in this way, then a check of access rights will be applied when the search is run even though the setting from the User, Group, or System object indicates otherwise.

Setting the Access Domain

The library system returns search results based on the user's access rights to the returned items. These access rights may be assigned:

- Directly to the user.
- To one or more of the groups to which the user belongs.

In the case of group access rights, the Access Domain property controls whether the library system returns search results based on the access rights of the user's active group or on the highest access rights of any of the user's groups.

Like most inheritance path properties in CS, an Access Domain property appears in User, Group, and System objects. Refer to [“Understanding Default Values and the Inheritance Path” on page 188](#) for details on how values are applied from these objects. The value in the User and Group object can be one of three values:

- Active Group
- All Groups
- Blank

The System object takes the same values, except that a blank value is not allowed.

Note Some interfaces may be designed to allow users to specify an access domain value when they create the search. If so, that value will override any default setting.

Enabling Content Searches

You can enable and disable the content search feature by toggling the Content Search Enabled? property value in the System object. You may want to disable content searches when doing certain kinds of system maintenance, such as restoring the Content Search Manager or manually rebuilding content search repositories.

Minimum Item Delete Access Level

Administrators can control the access level required to delete item objects by setting the System object's Minimum Item Delete Access Level property to either Admin or Owner. If set to Admin (the default), only library system administrators can delete items.

Refer to *CS Explorer Online Help* for information on how to edit the System object's Minimum Item Delete Access Level property.

Note This is a change in behavior from earlier versions of CS, in which item Owners were allowed to delete items.

If you change the value of the Minimum Item Delete Access Level property, active library system users must log off and log on again for this change to take effect.

Customizing Logon Procedures

When a library system is initially installed, it is configured so that users of IDM Desktop, CS Explorer, or other client applications must provide their library system user name and password (if any) in order to open a session. There are two ways the logon feature can be customized:

- You can restrict logons to administrators only, to prevent users from accessing the library system while you perform administrative tasks.
- You can configure your library system to allow the logon process to be bypassed altogether, so that a user's network logon name and password are used instead.

Restricting Logons to Administrators Only

As a library system administrator, you will sometimes need to prevent users from opening sessions while you are working in the System object or otherwise doing maintenance on the library system. Content Services provides the System object's Logons Enabled? property for this purpose. At installation, the value of Logons Enabled? is set to Yes. When you set the value to No, you are prohibiting any users from opening sessions until you set the value back to Yes. The only exceptions are for administrators (whose default group is set to Administrators), and the user (by default) named Admin, who can still open sessions even when Logons Enabled? has a value of No.

Refer to *CS Explorer Online Help* for information on how to edit the System object Logons Enabled property.

Using Network Logons

To enable a user to access a library system using network logon and password, the User object's Network Logon Name property must be set to *NetworkDomain;LogonName* for Windows NT Domain/2000 Active Directory directories. The value of this property is automatically set for users created by using any of the Import utilities. Refer to “Using Windows User Import” on page 273 for details.

The value assigned to this property must uniquely identify the user on the network. For example, if your network is configured to allow user *jdancer* to log onto the ENGINEERING domain, then setting the library-system User object's Network Logon Name property to ENGINEERING;jdancer enables that user to access the library system using network logon and password.

Also, your client application must enable the use of OS Credentials. Refer to the IDM Desktop/Web Services 4.x documentation for details.

Note this does not apply to LDAP Directories. To enable LDAP authentication, refer to the IDM Web Services documentation.

For details about using the Network Logon option in conjunction with the Windows User Import and LDAP User Import tools, refer to “[Passwords for Imported Users](#)” on page 284.

For information on how to edit User object properties (including the Network Logon Name property and the User Name property) refer to *CS Explorer Online Help*.

Note Users added to a library system via the User Import tool must either use the Network Logon option, or you as system administrator must change their password immediately after import. Refer to “[Passwords for Imported Users](#)” on page 284 for more information.

Assigning a Display Language

The Display Language property controls the user's display language for property names, field names, and messages (posted to the User object or the Administrators Group object). As with other inheritance path properties, the Display Language property appears in User, Group, and System objects. For details on how to set this value, as well as other values affected by the same inheritance path method, refer to “[Understanding Default Values and the Inheritance Path](#)” on page 188.

Note This Display Language property is usable only if your chosen user interfaces to the library system supports multiple language displays and if you have used a library system installation program to install the language data files required to support a display language other than those currently supplied. This kind of language support is available only with direct FileNet supervision. Contact FileNet Technical Support for details.

Changing Operating System and Database Display Language

Note Refer to “[Installing Content Services into Non-English Environments](#)” on page 81 before changing your operating system or database display language.

Use the following procedure to change the display language of your operating system and database (for example, from English to French). If you are using Replication Services, make sure that all servers for the participating library systems use the same language.

To change your operating system and database language:

1. If you are performing this procedure during an upgrade, first upgrade your Property Manager, refer to “[Upgrading Content Services](#)” on page 152.
2. Perform a complete backup of your library system. Refer to “[Backing Up the Library System](#)” on page 336.
3. Install Microsoft Windows 2000 or Windows 2003 and the service pack for the new display language.

4. Restore the library system property server. Refer to “Restoring the Database” on page 339 (note that while carrying out this procedure, you should restore the database engine in the language you want to change to.) Once you have restored the property server, return to this procedure.
5. For storage servers, repeat this procedure and refer to “Restoring Library Services” on page 340.

Controlling Library System Services

There are several ways to control library system services:

- The CS Service Manager, installed on each server during Content Services setup, allows you to start, stop, and view the status of library system services. Use CS Service Manager only if Content Services is installed in a *non-cluster* environment. Refer to “Content Services Manager” on page 198.
- The Windows Control Panel Services Tool provides control over the database services as well as the library system services. Refer to “The Control Panel Services Tool” on page 201.
- Using the command-line utilities installed with Content Services, you can start, stop, and view the status of library system services from a DOS command prompt. Refer to “Controlling Library System Services with Command-line Utilities” on page 202.
- If you have installed Content Services in a cluster environment, you *must* use Microsoft Cluster Administrator to start, stop, and view the status of library system services. Refer to “Controlling Library System Services in a Cluster Environment” on page 202

Content Services Manager

The CS Service Manager, which is automatically installed on the initial storage server and any additional storage servers during Content Services Setup, provides GUI controls and status display for starting, stopping, and monitoring library system services in a non-cluster environment.

Note If you are logged on to the storage server as a non-administrative Windows user, the CS Service Manager will not indicate that the Content Search Manager is running, even when it is running.

By default, the CS Service Manager is available as an icon in the System Tray at the far right of the Windows Taskbar of the server where you have installed library services. The icon becomes visible for the first time after you restart the server following the installation of the first set of library services on the server.

Whenever you finish installing another set of library services on the server, the CS Service Manager is not visible until after you have restarted the server.

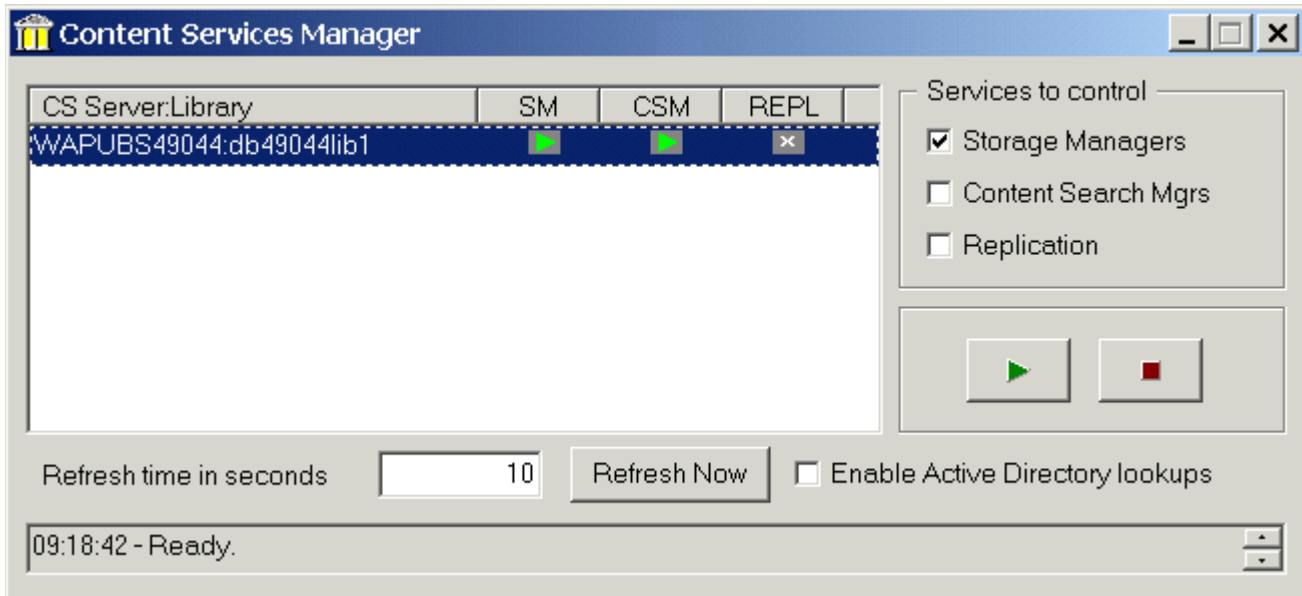
If you shutdown CS Service Manager, you can restart it by clicking Start > Programs > Startup > CS Service Manager, or by running *Drive:\filenet\dsstatus.exe*.

Content Services Manager without Active Directory Support

For Windows 2000 or Windows 2003 based library systems that have not been configured for Active Directory support, Content Services Manager can monitor and control local library system services only.

To use the CS Services Manager:

1. Double-click the CS Services Manager icon in the System Tray. The CS Services Manager window is displayed, showing the status of each library system component installed on that server.



In this example, the Storage Manager and the Content Search Manager for the *WAPUBS49044:db49044lib1* library system are currently running. Components showing a status of “X” (Replication Services in this example) are not installed.

2. To control services:
 - a. Select (highlight) the library system whose services you want to control. Or, to control the services of more than one library system at a time, hold down the shift key and select multiple library systems.
 - b. Select one or more check boxes to specify the type of service (or services) you want to control. In the above example, only the Storage Manager for *WAPUBS49044:db49044lib1* library system is selected for control.
 - c. Click the button for the action you want to perform.



Starts the selected service or services for all selected library systems.



Stops the selected service or services for all selected library systems.

The Service Manager queues requests to start or stop services. If multiple start or stop requests are made, they are executed in the order received.

Note The message text box at the bottom of the dialog box has scrolling arrows, enabling you to view the last 100 actions performed by the Service Manager.

3. To see the results of the action you performed, click Refresh. Each time you click Refresh Now, Content Services Manager will display the current status of the library system services. The “Refresh time in seconds”

text box indicates how frequently (in seconds) Content Services Manager will automatically update the status information of each displayed service.

Content Services Manager with Windows 2000 or Windows 2003 Active Directory Support

For library systems installed on Windows 2000 or Windows 2003 servers, the Content Services Manager can use Active Directory to enable you to monitor and control library system services on other servers within the same Active Directory forest.

Note If a library system service no longer exists on a server in the domain, use the Windows 2000 or Windows 2003 ADSI Edit MMC Plug-in (from the Windows 2000 or Windows 2003 Support Tools) to delete the service from the Active Directory list of services on that server. Use caution when deleting Active Directory objects with ADSI Edit, as incorrect modifications may adversely affect your entire domain tree.

To use the CS Service Manager with Active Directory Support:

1. Install and configure all Windows 2000 or Windows 2003-based library system servers within your Active Directory forest, as explained in: “[Installing Initial Library System Services](#)” on page 123.

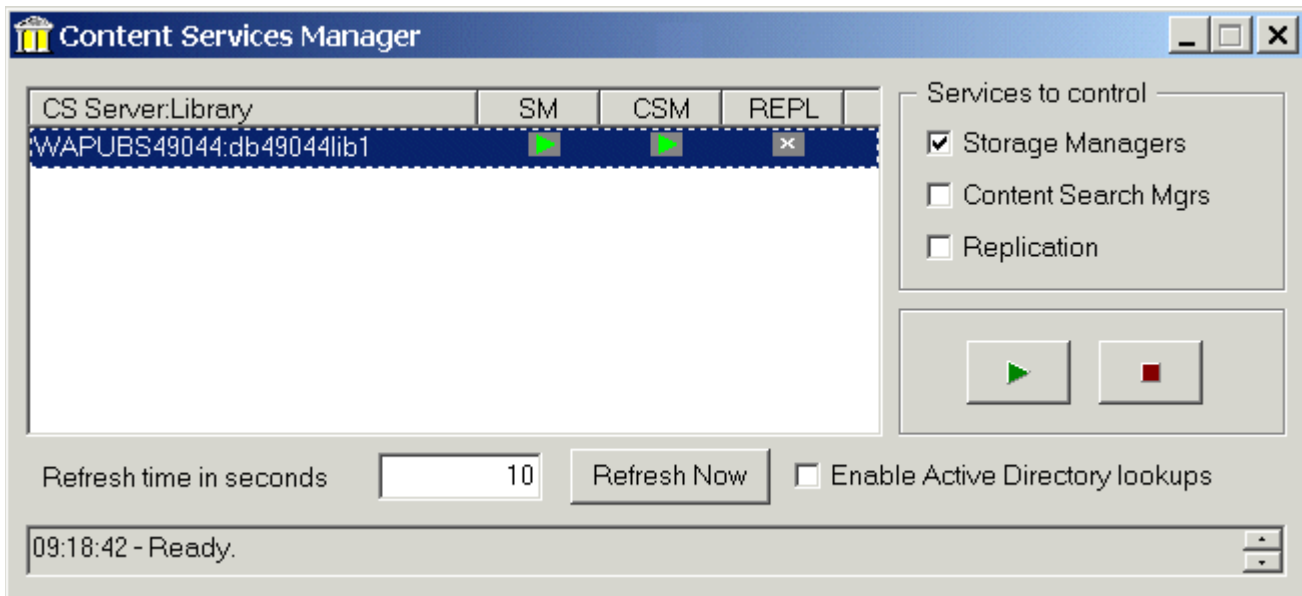
Note To enable remote monitoring and control of library system services, the Domain Admins and Enterprise Admins groups must belong to the local Administrators group on your library system server.

2. Log on to the library system server that you want to use for centralized control of your CS systems. The rights assigned to your User account will determine what library systems CS Service Manager can monitor and control:

| Members of | Can monitor and control |
|----------------------|---|
| Local Administrators | Local library system services only. |
| Domain Admins | Library system services on all Windows 2000 or Windows 2003 servers in the Active Directory domain. |
| Enterprise Admins | Library system services on all Windows 2000 or Windows 2003 servers in the Active Directory forest. |

3. Double-click the Content Services Manager icon in the System Tray. The Content Services Manager window is displayed, showing the status of each library system component installed on that server.

In this example, the Storage Manager and the Content Search Manager services for the *WAPUBS49044:db49044lib1* library system are currently running. Components showing a status of “X” (Replication Services in this example) are not installed. If the number of library systems exceeds the available space in the Service Manager window, a vertical scroll bar appears to allow you to scroll up and down the list.



4. You control library system services on remote servers just as you would on the local server (see [“Content Services Manager without Active Directory Support”](#) on page 198) above. Select or clear the check “Enable Active Directory lookups” box as needed.

The Control Panel Services Tool

You can use the Windows Control Panel Services tool to start, stop, pause, or continue library system services and your database services.

Note The Windows Oracle edition of Content Services does not support automatic startup of library system services.

To control the services using the Windows Control Panel Services tool:

1. Access the Windows Services tool in Control Panel > Administrative Tools.
2. Double-click the **Services** tool to start it.
3. If you are going to start library system services, first verify that the your database services are started.
 - MSSQLServer (for Windows-SQL)
 - OracleServiceORACLE_SID (for Windows-Oracle)
 - OracleStartORACLE_SID (for Windows-Oracle)
 - OracleTNSListenerLISTENER_NAME (for Windows-Oracle)
4. Locate the services for your library system. If Content Services is correctly installed, you will see a service name for each component:
 - Storage Manager - *LibrarySystem* (required)

- Content Search Manager - *LibrarySystem* (optional)
 - Replication Manager - *LibrarySystem* (optional)
5. Select the service you want to control. Right-click the service and then select Start, Stop, Pause, Resume, or Restart from the context menu to put the selected service into the desired state.
 6. For Microsoft SQL Server systems, you can set your library services to start up automatically. Refer to your Microsoft document for instructions.

Controlling Library System Services with Command-line Utilities

Command-line utilities for starting, stopping, and checking the status of library system services are installed on each library system server. To run these utilities, start a command prompt and change to the *Drive:\filenet* directory, then enter the command.

| Utility | Description |
|---------|--|
| DSINFO | Displays the status of library system services, as well as the Property Manager (database engine), for the specified library system. By default DSINFO displays status information for all library system services on the machine where it is run. Refer to “DSINFO” on page 398 for complete documentation. |
| DSSTRT | Starts library system services for the specified library system. By default, DSSTRT starts all library system services (except Replication) for all library systems on the machine where it is run. Refer to “DSSTRT” on page 402 for complete documentation. |
| DSSTOP | Stops library system services for the specified library system. By default, DSSTOP stops all library system services on the machine where it is run. Refer to “DSSTOP” on page 400 for complete documentation. |

Controlling Library System Services in a Cluster Environment

When you are running Content Services in a cluster environment, you must use Cluster Administrator to toggle the state of a library system between online and offline. Do *not* use the CS Service Manager, the Control Panel Services Tool, or command-line utilities to toggle the state.

To bring a library system online:

1. Start Cluster Administrator. The Open Connection to Cluster dialog box opens (same as choosing File > Open Connection). Select from the dropdown list the cluster name where Content Services is installed and click Open.
2. On the left pane, expand the Groups folder icon. Then select the CS_*libraryname* icon, where *libraryname* is the name of the library system. You can either bring all the services for a library online or just an individual service.
 - a. If you want to bring all the services for the library online, right-click the CS_*libraryname* icon in the left pane and choose Bring Online from the pop-up menu.

- b. If you want to bring an individual service online, right-click the CSMGR_*libraryname*, REPL_*libraryname*, or SHF_*libraryname* icon in the right pane and choose Bring Online from the pop-up menu.

Note If at least one service for a library system is not running, the CS_*libraryname* icon in the left window will display a red X.

After bringing your library systems online for the first time, continue at “[Accessing and Administering a Library System](#)” on page 130.

To take a library system offline:

1. Start Cluster Administrator. The Open Connection to Cluster dialog box opens automatically (same as choosing File > Open Connection). Select from the dropdown list the cluster name where Content Services is installed and click Open.
2. On the left pane, expand the Groups folder icon. Then select the CS_*libraryname* icon, where *libraryname* is the name of the library system. You can either take all the services for a library offline or just an individual service.
 - a. If you want to take all the services offline, right-click the CS_*libraryname* icon in the left pane and choose Take Offline from the pop-up menu.
 - b. If you want to take an individual service offline, right-click the CSMGR_*libraryname*, REPL_*libraryname*, or SHF_*libraryname* icon in the right pane and choose Take Offline from the pop-up menu.

Accessing Multiple Library Systems

Content Services provides flexibility in how you can configure your library systems and how your users can access them through their CS client applications.

Refer to:

- [“Multiple Library Systems With One Database Engine” on page 204](#) for information about installing multiple library systems on a single property server, using a single database engine
- [“Connecting a Client to Different Library System Platforms” on page 206](#) for information about allowing clients to communicate with library systems running on different servers with different database platforms

Note (FNDTS00071308) A network logon to a library system from a client application will fail if you try to log on directly to the Windows domain after canceling a Novell Client for Windows logon. To avoid the logon failure, perform the following steps on your client workstation before trying to log on:

1. Create the user environment variable **sni_network_type** and set its value to 1.
2. Reboot your workstation.

Multiple Library Systems With One Database Engine

Each library system has an associated property repository. The property repository is a relational database that maintains information about the library system, including the network configuration, the library system's users, and the files and other items that it manages (including the status, network location, and general descriptive properties for each item).

The database engine (for example, Windows Oracle) is typically used by a single library system. However, the database engine can be shared by several library systems.

We recommend that no more than four library systems share a database on a single CPU. Additionally, the minimum memory requirements for up to four library systems are:

| Number of Library Systems | Memory Required |
|---------------------------|-----------------|
| 1 | 160 MB |
| 2 | 240 MB |
| 3 | 320 MB |
| 4 | 400 MB |

Note More than four library systems would require an SMP server or a different server platform.

FileNet certifies Content Services for use with up to twelve library systems sharing a database engine. Refer to [“Configuring Databases” on page 229](#) for instructions on reconfiguring the library systems' database parameters to assign more memory to the database.

This type of multisystem configuration requires that the server on which this software is installed be able to handle this kind of load. It also assumes that you would want to locate the property repository for multiple library systems on a single CPU — a configuration that can be economical of hardware resources, but can sometimes complicate maintenance, backup/recovery, and even software update procedures. For example, if you have hardware failure on a server containing multiple library systems, more users are likely to be adversely affected than if you had spread the library systems over multiple servers so that some library systems could still be functioning. Recovery time is also going to be lengthened if you have to restore more than one library system per server.

Note You can always install multiple library systems at your site, each with its own database engine. This configuration tends to be simpler to maintain, and these totally independent library systems can still be accessed from a single client with little trouble and minimal setup requirements.

Before You Install Additional Library Systems

There is no simple way to ensure in advance that the server where the property repository is to be installed for your various library systems has adequate resources. The actual requirements for your site depend on the operations being performed on the additional library systems, including how many users will be working concurrently and how many files will be checked in on a daily basis. Review the requirements in [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#).

Refer to “[Managing Multiple Library Systems](#)” on page 310 for details about configuring more than one library system at your site.

Adding a Library System

The following general guidelines assume that you are installing (not upgrading) an additional library system onto a server that already has a fully upgraded or newly-installed library system. The operating system and database software therefore already exists on servers and clients. For hardware and software requirements, refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#).

1. Refer to “[Before You Install Additional Library Systems](#)” on page 205 and verify that you have completed the necessary steps.
2. In the existing database, create a tablespace for the new library system. Refer to “[Creating a Property Repository Tablespace](#)” on page 106.

The new library system will use your existing ODBC data source (DSN).

3. Install the new library system. Refer to “[Installing Content Services](#)” on page 119 for installation instructions. Install additional Storage and Content Search Managers as required.

Note Installing an additional library to a property server shuts down the existing libraries. You will have to restart them manually.

4. Configure your CS Explorer to access the new library system.
5. Configure each client application on each client machine that will access the new library system.
6. Refer to “[Replication Services](#)” on page 593 for instructions if the new library system will be part of an existing replication domain.

Connecting a Client to Different Library System Platforms

A single CS client application can access multiple library systems, each supported by a different database engine. For example, you could set up a single CS client workstation to connect to a library system on one server, supported by a Windows Oracle database, and to connect to another library system on another server, supported by a Solaris Oracle database, and to yet another library system on a third server, supported by a Microsoft SQL Server database.

Refer to “[Client Library Files](#)” on page 207 for information on how to enable access from a CS client application to a library system supported by an Oracle database. Refer to the Windows SQL edition of the Content Services documentation at [FileNet Worldwide Customer Support](#) or information on how to enable access from a CS client application to a library system supported by Microsoft SQL Server database engine.

Client Library Files

The Content Services Client Library Files, included with all Content Services client distributions (such as IDM Desktop), are the files that enable a CS client application to access a CS library system. The CS Client Library Files include the client files required by the library system, system DLLs such as **msvcr7.dll**, and ODBC files. The Oracle Client files and an ODBC Data Source Name are also required on the machine where your CS client application is to be installed. Refer to [“Installing the Oracle9i Database Client Files”](#) on page 109 or [“Installing the Oracle 10g Database Client Files”](#) on page 111 for instructions.

In most cases, you will not need to install CS Client Library Files as a separate procedure. In this scenario, the CS client application's installation program calls the Client Library Files Setup program to run in the background. In addition to installing the necessary libraries, Client Library Files Setup:

- Sets path statements
- Enters information in the Windows registry
- Tests (on next restart) the connection to the library system database

However, independent installation of the CS Client Library Files is sometimes required because this software has a product lifecycle of its own and can be distributed independently. An independent installation may be required to upgrade an existing client installation, replace a damaged installation, or complete the installation of a non-FileNet client application that does not include the CS Client Library Files. Refer to the CS client application documentation if you are installing the CS Client Library Files as part of a client application installation.

The following topics describe how to enable access from a CS client application to a library system supported by an Oracle database engine. But you can also enable access from a CS client application to a library system supported by Microsoft SQL Server database engine. That is, you can enable a CS client application to connect to multiple library systems, each based on a different database engine: Microsoft SQL Server, HP Oracle, Windows Oracle, or Solaris Oracle.

Refer to the Windows SQL edition of the Content Services documentation at [FileNet Worldwide Customer Support](#) for information on how to enable access from a CS client application to a library system supported by Microsoft SQL Server database engine.

Prerequisites to Installing Client Library Files

Before starting an installation of the CS Client Library Files, refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#) to make sure that your client machine complies with hardware and software requirements. Then do the following:

1. Install the Oracle Client software on your client machine. Refer to [“Installing the Oracle9i Database Client Files”](#) on page 109.
2. Establish connectivity between your client machine and the Oracle database used by the CS library system your CS client application will be accessing. Refer to [“Preparing Oracle9i or Oracle 10g Client Connectivity”](#) on page 112.
3. It is recommended that you uninstall any existing CS 5.4 Client Library Files if you do not need to perform content-based searching against pre-CS 5.5 library systems. If you do not uninstall the files now, and you want

to uninstall them later, you will need to reinstall the CS 5.5 Client Library Files after uninstalling the CS 5.4 Client Library Files.

Installing Client Library Files

To install the CS Client Library Files on a Windows machine, you can either run the interactive Client Library Files Setup program, as described in this topic, or you can install the files unattended (noninteractively), as described in [“Installing Client Library Files Unattended” on page 208](#). To install the CS Client Library Files on a UNIX machine, refer to [“All UNIX Clients” on page 214](#).

To install the Client Library Files interactively:

1. Insert the *FileNet Content Services 5.5 Admin Tools* CD-ROM in the CD-ROM drive of the client machine and double-click **Drive:\Client.Lib\Win32\Language\Setup.exe** on the *Content Services Admin Tools* CD-ROM.
2. In the Choose Destination Location dialog box, click Next to accept the default (or click Browse if you want to specify a different target directory).

Note If a previous installation of the Client Library Files is detected, Setup will use that location for the default.

3. Setup will install and configure the Client Library Files. In the Setup Complete dialog box, you can choose to reboot your system now or later. Click Finish to exit.
4. Check the output file for information about the installation. By default, Setup writes an output file to **Drive:\Temp\Setupout.ini**. Refer to [“Client Library Files Setup Output File” on page 211](#).

Installing Client Library Files Unattended

You can run Client Library Files Setup from a command line (using the a **Clissetup.ini** file and a **Setup.iss** file) to perform *unattended* or *silent* client installation.

Preparing a Clissetup.ini File

To perform unattended Client Library Files installation, you must first prepare an initialization file to provide input to Setup. You can prepare an initialization file either by editing the provided **Clissetup.ini**, or by running Setup with the `-ADMCFG:filename` parameter (refer to [“Client Library Files Setup Command-line Options” on page 209](#)).

To edit the provided Clissetup.ini file:

1. Copy **Drive:\Client.Lib\Language\Clissetup.ini** from the *Content Services Admin Tools* CD-ROM to a location accessible to the machine on which you will run Setup.
2. Edit **Clissetup.ini** with a text editor. Refer to [“Client Library Files Setup Input File” on page 210](#) for format and content requirements.
3. You can now use your initialization file to install Client Library Files.

Note You can name the initialization file anything you wish. You will specify the file to use by including `-IN:filename` on the Setup command line.

Performing a Silent Installation of Client Library Files

Once you have prepared a **Clissetup.ini** file, as explained above, you are ready to run Setup to perform a client installation.

To install Client Library Files using Clisetup.ini:

1. Start a DOS command prompt and enter a command to perform the installation. For example:

```
setup -IN:Clisetup.ini -s -f1 Setup.iss
```

This command will install the client libraries according to the setup parameters in Clisetup.ini. Command prompts will not be displayed. Refer to “Client Library Files Setup Command-line Options” on page 209.

2. Check the output file for information about the installation. By default, setup writes an output file to **Drive:\Temp\Setupout.ini**. Refer to “Client Library Files Setup Output File” on page 211.

Client Library Files Setup Command-line Options

The Client Library Files Setup program supports the following command-line options. Options can appear in any order on the command line, provided the InstallShield options appear after all other options. Avoid using spaces in filenames.

| Option | Required/Optional | Description |
|--------------------------|---|---|
| -IN: <i>filename</i> | Optional in UI mode; Required in non-UI mode | Tells Setup to perform the installation using the installation parameters contained in <i>filename</i> . Refer to “Client Library Files Setup Input File” on page 210. |
| -OUT:[<i>filename</i>] | Optional | Tells Setup to write information about the installation to <i>filename</i> . If you omit this option, Setup writes the output to Drive:\Temp\Setupout.ini . Refer to “Client Library Files Setup Output File” on page 211. Note If the -OUT: option is used without a filename, output is suppressed. |
| -LOCATION: <i>dir</i> | Optional | Target directory for the client libraries installation. If you omit this option, the default value is used. For Web installations the default is the Windows System folder. For Desktop installations, the default is Drive:\Program Files\FileNet\Shared . If the client libraries have previously been installed, this value is ignored and the location of the previous installation is used. |
| -AGGREGATE: | Optional | Tells Setup that the client libraries installation is a part of another process. Welcome and Finish dialogs are suppressed. This option has no effect in silent mode. |

| Option | Required/Optional | Description |
|---------------------|-------------------|---|
| -SMS | Optional | (Microsoft SQL Server databases only) Provides support for installations via Microsoft Systems Management Server (SMS). Specifically, this flag causes the installer to spawn processes synchronously rather than asynchronously, so when the parent process (task) terminates, SMS can be assured that the entire installation has finished. |
| -s | Optional | Specifies silent mode. Refer to Microsoft InstallShield documentation for details. |
| -f1 <i>filename</i> | Optional | Calls <i>filename</i> (a setup file) to be used during silent mode installation. A sample setup file is provided in Drive:\Client.Lib\ Language\Setup.iss on the <i>Content Services Admin Tools</i> CD. Refer to Microsoft InstallShield documentation for details. |

Usage examples

Note Flags in bold are required, and the order is significant. The *.iss file is required by InstallShield to run in silent (non-UI) mode. Refer to the InstallShield documentation for details.

| You want to... | Command-line example |
|--|---|
| Install in UI mode, to let the user control the installation process | <code>setup -IN:params.ini</code> |
| Run the installation as a part of another process, in non-UI mode | <code>setup.exe -AGGREGATE: -IN:myfile.ini -s -f1 setup.iss</code> |

Client Library Files Setup Input File

Client Library Files Setup can read an input file if it is specified on the command-line. The Input file provides a way for the calling application, which could be a program you created yourself, to configure the installation and run it in batch mode without user interaction. This is called unattended or silent mode.

The Input file is a standard Windows initialization (*.ini) file, which can be located in any read-accessible location. Use the command-line parameter -IN to specify the exact location. For example:

```
setup.exe -IN: path\infile.ini -LOCATION: target-path -s -f1 path\ setup.iss
```

The Input file includes the following:

| Entry | Values | Description |
|--------------------|--------|---|
| InstallFileVersion | 2 | Required. Indicates the version of the input file. Due to some significant changes in the file structure, the installer version 2 doesn't accept files created for installer version 1. |

| Entry | Values | Description |
|-----------------|----------------|---|
| SkipMDAC | 0 or 1 | <p>Optional. Default is 0. If set to 1, the MDAC installer will not be called. Note that the installer will skip execution of MDAC automatically if current or newer MDAC components already exist on the system. It is only necessary to use this option if you want to force MDAC to be skipped.</p> <p>If set to 1 and ODBC does not exist, the Client Library Files Setup will fail.</p> |
| TargetDirectory | Any legal path | <p>Optional. The target installation directory for the Common Files group.</p> <p>TargetDirectory takes precedence over the -LOCATION:<i>dir</i> command-line parameter.</p> <p>To prevent multiple installations of the Client Library Files in different locations on the same machine (which can cause DLL conflicts), TargetDirectory is ignored if the following registry key exists:</p> <pre>HKEY_LOCAL_MACHINE\SOFTWARE\FileNET\Pa nagon Content Services Client Libraries\version\InstallPath</pre> <p>This registry key indicates that Client Library Files have been installed at least once. The value of InstallPath will then be used as the target directory for any subsequent Client Library File installations.</p> |

Note The application installer that supplied the input file is responsible for cleaning up any temporary or intermediate input files when the Client Library Files Setup completes.

Client Library Files Setup Output File

The Client Library Files setup program always generates an output text file, unless you suppress output by using the **-OUT:** command-line option without a filename. The default location where the output is written is **Drive:\Temp\Setupout.ini**. Refer to “Client Library Files Setup Command-line Options” on page 209 for more information.

Note Refer to “Client Library Files Setup Error Messages” on page 212 for a list of the potential error codes and messages that could be reported in the output file.

An output file can have the following entries.

| Section | Entry | Description |
|------------------------|---------------------------------|--|
| [Result] | ReturnCode | Completion Code: 0 = OK 2 = Essential steps OK, but some problems occurred 4 = A fatal error has occurred 12 = Cancelled by user (incomplete) |
| | Warnings | Might not appear. A number of warnings are possible. |
| | ErrorCode | Might not appear. An error code for fatal error (if a fatal error occurred). |
| | errorcode | Might not appear. Error message for error above. |
| [Warning- <i>nmn</i>] | MustReboot | 0=User may reboot later 1=Some important functionality will take effect only after reboot |
| | ErrorCode | Might not appear. An error code for each warning. |
| [Path Conflicts] | errorcode | Might not appear. Warning message for the code above. |
| | TotalPaths= <i>n</i> | The total number of paths where path conflicts (described below) were detected. |
| | Path <i>n</i> = <i>pathname</i> | Path <i>n</i> numbering begins with Path1. The pathnames listed here are paths where srs3mapi.dll, srs3dbcf.dll, or srs3odcf.dll were detected in locations <i>other than</i> the location where setup was trying to install the Client Library Files. |

Client Library Files Setup Error Messages

The following is a list of the error codes and messages that may be reported during the Client Library Files installation process. Note that some errors might be fatal or non-fatal depending on the context. For more library system error codes, see “[Status Codes](#)” on page 573.

| Error code | Message and meaning |
|------------|--|
| 5000 | The “5000” code means “General Error” and most likely could be caused by an internal InstallShield problem, such as media corruption or memory leak. Rare. |
| 5001 | An error occurred during the data move process. Rare. |
| 5002 | Internal component(s) error. Most likely that you have a broken installation. Rare. |
| 5010 | Install couldn't create an internal list. Most likely, you ran out of memory. Very rare. |

| Error code | Message and meaning |
|------------|--|
| 5012 | Error while adding/updating an internal list. Very rare. |
| 5020 | Install couldn't load external DLL. Most likely, the installation was corrupted. Rare. |
| 5030 | This program requires Windows 98 or higher. |
| 5040 | Invalid command line. |
| 5042 | You cannot use these command line options together |
| 5050 | Could not create or append the file. |
| 5055 | Cannot access the directory for write/append |
| 5056 | Could not create the directory |
| 5060 | Could not find the file |
| 5061 | Couldn't find the command file - the file specified by the -INCMD flag doesn't exist or cannot be accessed. |
| 5062 | Directory doesn't exist |
| 5063 | Sqlnet.ora was not found. |
| 5065 | The target folder doesn't have enough space |
| 5070 | Uninstaller setup failed to initialize. You may not be able to uninstall this product. In some cases might be caused by memory or registry related problems. Rare. |
| 5080 | No supported database software has been found on your system. |
| 5090 | Could not initialize some custom dialog box resources. Most likely, the installation was corrupted. Rare. |
| 5092 | Internal error while showing a custom dialog box. Very rare. |
| 5100 | Error accessing Registry |
| 5105 | Error updating Registry |
| 5110 | Invalid input value. Some specific entry in the input file is invalid. |
| 5111 | Invalid Input file version - the input file doesn't have version key or the version is less than 2. |
| 5150 | Error while getting product version - Not used in this release. |
| 5170 | Couldn't complete self-registration process |
| 5190 | Please be advised that in order to perform this installation you must have Administrator privileges. Windows 2000 or Windows 2003 only. |

| Error code | Message and meaning |
|------------|--|
| 5200 | Error accessing environment - Installer couldn't get some environment variable(s) settings. Very rare. |
| 5205 | Error updating environment. Very rare. |
| 5300 | MDAC Installer failed. |
| 5302 | DCOM Not Found (Windows 98 only). |
| 5303 | One or more sets of conflicting Client Library Files were found via the PATH. |
| 5304 | Read-only files were found (silent-mode warning). |
| 5307 | DBMS client libraries were upgraded (silent-mode warning). |
| 5308 | Install path overridden by registry stored path. Client Libraries being installed in: <i>dir</i> . |

All UNIX Clients

The shared Client Library Files for UNIX clients have their own install program, which can allow for easier intermediate updates of UNIX client libraries without requiring recompiles of the entire UNIX client application. You can run the install program interactively to update client libraries on particular UNIX clients (say, for third-party client applications).

Before you install the CS Client Library Files on the UNIX client machine, make sure you have installed the Oracle Client files corresponding to the Oracle version (Oracle9i or Oracle 10g) of the database used by Content Services. Refer to “[Connecting to the Oracle Database](#)” on page 109.

The following installation procedure assumes that a UNIX user is installing the CS Client Library Files into the directory **\$HOME/CSlibs**.

To install the shared Client Library Files on a UNIX machine:

1. Make sure that you have installed the Oracle Client Library files.
2. Determine/create an installation directory on the UNIX client.
3. Mount the *Content Services Admin Tools* CD-ROM.
4. On the CD-ROM, change to the subdirectory for your UNIX (HP-UX or Solaris) platform:


```
cd /Client.Lib/unix_platform
```

 where *unix_platform* is HPUX or SOL.
5. Type:


```
./install
```
6. Answer the prompts for:
 - a. Path to which to copy the shared libraries (**\$HOME/CSlibs**).

Problem Solving

If the installation/upgrade of the library system services appears to complete successfully but the library system will not run, review the following log files to check for error messages:

| Log | Description |
|--------------------------------|---|
| cssetup.log | A log (located at <i>InstallDrive:\filenet</i> by default) of the individual actions (file copying/deleting, component installation, custom property creation, etc.) in the installation of the initial/additional library services. |
| GUISetup- <i>username</i> .log | A log file (located at %TEMP% (where %TEMP% is an environment variable for <i>username</i> , the user installing Content Services on this server) to see a log of the screen-related actions of CS Setup in the installation of library services. |

To obtain the meaning of a Verity-related error message in the log file:

1. At a DOS prompt, invoke the Verity command **rcadmin**.
2. At the rcadmin prompt, type **errormessage**.
3. At the error-code prompt, type the number of the error; **rcadmin** returns a description of the error.

The following problems may also occur while you are setting up the library system services. If a problem does occur, follow the recommended action to solve it.

| Problem | Possible Cause/Solution |
|-------------------------|---|
| Insufficient disk space | Check the disk-space requirements for the library services in CS 5.5.0 Windows Oracle Platform Requirements at FileNet Worldwide Customer Support . If an error message indicates inadequate disk space, you must either remove old files or install the library system services on a disk with more available space. |
| Insufficient memory | This error message indicates that the library system host has no more available RAM. If you receive an error message indicating a lack of RAM, you should: <ul style="list-style-type: none"> • Check to see if other applications on the server are running. Make sure they are in an idle state or, preferably, shutdown before trying to install the library system again. • Check to see that database engine is currently configured for the correct amount of RAM. • Check to see that all the available RAM is accessible and error free. • Add more memory. |

| Problem | Possible Cause/Solution |
|--------------------------------------|---|
| Cannot log in to database | <p>This error message indicates that the Storage Manager cannot communicate with the installed database. If you receive this error message, do the following:</p> <ul style="list-style-type: none"> • Make sure that you have properly installed the network adapter card and any required drivers. • Run the PING utility to ensure that the network communications are functioning properly. • Make sure the name for any server on which library system software is installed is included in the %SystemRoot%\system32\drivers\etc\hosts file. This file will also include the IP address for each server. • Make sure you have created an ODBC connection. • Check whether you can still log on to the property server over the network using a database client tool (SQL*Plus or Oracle Enterprise Manager, for example). |
| Power interruption | <p>If power to the storage server is interrupted during installation, start the library system installation from the beginning.</p> |
| Storage Manager fails to bind socket | <p>This error condition can result when you attempt to restart a Storage Manager soon after a shut down. If you receive this error (226), wait at least five minutes and then try to start the Storage Manager again</p> |

Uninstalling Content Services

To uninstall Content Services, you uninstall library system services and client library files. To uninstall library system services, refer to:

- [“Uninstalling Library Services” on page 217](#) to uninstall a library system using the Windows Control Panel Add/Remove Programs tool
- [“Uninstalling Library Services Unattended” on page 218](#) to uninstall a library system using the Dsuninst program
- [“Uninstalling Content Services in a Cluster Environment” on page 219](#) to uninstall a library system in a cluster environment.

However you decide to uninstall library services, when uninstalling a Windows 2000 or Windows 2003-based library system that is configured for Active Directory support, be sure to log on as a user with Domain Admin rights or higher. This enables the uninstall program to remove the library system's Active Directory objects from the domain controller. If you do not log on with at least Domain Admins rights, the Active Directory objects for the library system you are uninstalling will not be deleted from the domain controller and the library system will continue to be listed in the CS Service Manager with an “X” status indicator for all components. In that case, to remove all reference to the deleted library system from the CS Service Manager, you would have to use a special tool (such as Microsoft ADSI Edit) to manually delete the library system's Active Directory objects.

Uninstalling library services uninstalls the Content Search Manager, including the Verity software. But you might need to uninstall just the Verity software for a set of library services, or for all sets of library services, on a server. For example, the library services on the server may be referencing a Verity K2 Master Admin that no longer exists and now need to reference a new Master Admin. To uninstall the Verity software, refer to [“Uninstalling Verity Software” on page 218](#).

To uninstall client library files, refer to:

- [“Uninstalling Client Library Files” on page 219](#)
- [“Uninstalling Client Library Files Unattended” on page 220](#)
- [“Uninstalling Client Library Files via MS-SMS” on page 220](#)

Uninstalling Library Services

The following procedure shows how to remove one or more sets of library services from a server. Before uninstalling library services, if you want to retain the library system data stored in the database, use your database tools to do a full backup and save your backup files, the files from your storage repositories, your content search repositories, and your content search data files to a new location. Refer to your database documentation on backing up and to [“Backing Up the Library System” on page 336](#).

In Content Services 5.5, uninstalling a set of library services deletes the stacks directory and all documents stored in the directory. In some previous versions of Content Services, uninstalling left the stacks directory and its contents intact. Uninstalling a library system does not remove the property repository files.

To remove a set of library services:

1. On the server from which the library services are to be removed, open the Windows Control Panel and start the Add/Remove Programs tool.
2. While the Add/Remove dialog box is displayed, close the Control Panel window. This is necessary so that your `\winnt\system32\Dsproperties.cpl` can be deleted.
3. On the Add/Remove Programs screen, select **FileNet Content Services 5.5** and click Change/Remove. This launches the Content Services Setup program, which will guide you through the rest of the uninstall.
4. On the Remove Library System screen, select the library system whose services you want to uninstall from the server. Then click Next. (Click Yes when prompted to remove the library services.)
5. Note the results of the uninstall on the Remove Library System screen. Select the check box if you want to uninstall the services of another library system on the server; clear the check box if you do not want to uninstall any other library services. Then click Next. If you selected the check box, continue at [step 4](#).

Uninstalling Verity Software

Invoke the DSUNINST utility to remove the Verity software for a set of library services, or for all sets of library services, on a server. For example, the library services on the server may be referencing a Verity K2 Master Admin that no longer exists and now need to reference a new Master Admin.

WARNING Before running `dsuninst`, back up the indexes for each set of library services you have uninstalled. The `\content\index` directory will be removed during the Verity uninstall process.

1. Run `dsuninst` to uninstall the Verity software:
 - If at least one set of library services is to remain on the server, run `dsuninst` once for each set of library services whose Verity software you need to uninstall.


```
dsuninst -v library_name
```
 - To uninstall the Verity software for every set of library services on the server, run `dsuninst` once:


```
dsuninst -v
```
2. If you have removed the Verity software for every set of library services on the server, manually delete the `dsuninst` program and log file from the `\filenet` directory, and then remove the directory.

Uninstalling Library Services Unattended

As an alternative to running the interactive Content Services Setup program to uninstall library services on a server, you can run an unattended (silent noninteractive) uninstall by running the utility `dsuninst` at a command line.

To uninstall library system services:

1. Run the `dsuninst` utility on a command line once for each set of library services you want to uninstall on the server:

```
dsuninst -i LibrarySystem
```

Uninstalling Content Services in a Cluster Environment

For each set of library services you want to uninstall, perform the following procedure.

To uninstall Content Services:

1. Start Cluster Administrator on any node of the cluster and take the Content Services groups on all the cluster nodes offline. (See “Controlling Library System Services in a Cluster Environment” on page 202). This stops all library system services.
2. In the Add/Remove Programs applet of the Control Panel, select FileNet Content Services 5.5 and click Change/Remove.
3. At the Remove Library System screen of CS Setup, choose a library from the list of library systems and click Next.
4. Repeat step 3 if you wish to uninstall another set of library services on the server; otherwise click Cancel.
5. Power down the non-cluster owner node.
6. Reboot the cluster owner node and log on as a user in the Domain Admins group.
7. Restart the non-cluster owner node and log on as a user in the Domain Admins group.

If at this point any library services remain installed on the server, skip the remainder of this procedure.

8. The uninstall may fail to remove some Content Services-related keys in the Windows registry. Search the registry for any remaining keys (look for keys containing “IDMDS” or “Verity”) related to Content Services and remove them manually. Also, make sure to remove the FileNET key from the SOFTWARE key in the Registry.

Uninstalling Client Library Files

To uninstall FileNet's Client Library Files:

1. Open the Control Panel and start the Add/Remove Programs tool.
2. In the Install/Uninstall tab, select FileNet Content Services Client Libraries 5.x.
3. Click Add/Remove.

Uninstall Notes

- This uninstall procedure does not remove Microsoft ODBC, OLE DB, ActiveX Data Objects, and Remote Data Services files. Although these files were installed by Client Library Files Setup, they may be required by other applications that were installed after the Client Library Files. Refer to your database documentation for the correct uninstallation procedure, or contact Microsoft.
- If you uninstall a library system and then uninstall the Client Library Files (which is the correct order) the directory in which the Client Library Files were installed will still be present and will still contain several DLL files. You can delete these if you wish.

Uninstalling Client Library Files Unattended

You can run the command **csuninst32** on a command line to perform a silent (noninteractive) uninstall of the Content Services 5.5 or 5.4 client library files. This type of uninstall removes the client files on the machine to which you are logged on.

To silently uninstall Content Services 5.5 client library files, invoke the following command:

```
csuninst32 -v5.5 -s
```

To silently uninstall Content Services 5.4 client library files, invoke the following command:

```
csuninst32 -v5.4 -s
```

Uninstall Notes

Csuninst32 does not remove Microsoft ODBC, OLE DB, ActiveX Data Objects, and Remote Data Services files. Although these files were installed by Client Library Files Setup, they may be required by other applications that were installed after the Client Library Files. Refer to your database documentation for the correct uninstallation procedure, or contact Microsoft.

If you uninstall a library system and then uninstall the Client Library Files (which is the correct order) the directory in which the Client Library Files were installed will still be present and will still contain several DLL files. You can delete these if you wish.

Uninstalling Client Library Files via MS-SMS

You can use Microsoft Systems Management Server (SMS) to perform silent (noninteractive) uninstallation of the Client Library Files.

To uninstall Client Library Files using MS-SMS:

1. Using SMS, create a distribution package for the uninstallation. As a starting point for the content of your package, you can import *Drive:\Client.Lib\Language\Clisetup.pdf*. Refer to the Microsoft SMS documentation as necessary.
2. Choose the **ClientUninstall** command.
3. The command **csuninst32** performs a silent uninstall of the Content Services 5.4 or 5.5 client library files.

To uninstall Content Services 5.4 client library files:

```
csuninst32 -v5.4 -s
```

To uninstall CS 5.5 client library files:

```
csuninst32 -v5.5 -s
```

Place these commands in the corresponding `CommandLine` entries of **Clisetup.pdf**

4. Use SMS to create a job and schedule the uninstallation. Note that once the job is created, it may take quite awhile (an hour or more) for it to arrive at the computer where it is to be run.
5. On the computer where the library system is to be uninstalled, execute the job to run the distribution package. When the uninstallation completes, the SMS Administrator will show the status.

Administration Guide

| For information on: | Refer to: |
|---|---|
| An overview of library system administrative and maintenance tasks | “Configuration And Maintenance” on page 224 |
| Reconfiguring database parameters to improve system performance | “Configuring Databases” on page 229 |
| Adding, modifying, or moving the repositories where your documents are stored, or setting up direct-access storage repositories | “Storage Repositories” on page 233 |
| Adding or modifying categories or groupings of documents with common storage requirements | “Storage Categories” on page 240 |
| Moving item versions offline, reclaiming those versions when necessary and adding, modifying, or moving repositories where your documents are archived. | “Archive Repositories” on page 242 |
| Adding or modifying categories or groupings of documents with common archiving requirements | “Archive Categories” on page 250 |
| Adding, modifying, or moving the repositories where your content searches are stored | “Content Search Repositories” on page 253 |
| Adding or modifying custom properties, which can be used (in addition to the standard properties) to identify, locate and retrieve documents | “Custom Properties” on page 258 |
| Working with User objects | “User Objects” on page 263 |
| Working with Group objects | “Group Objects” on page 267 |
| Setting up and managing multiple library systems at your site | “Managing Multiple Library Systems” on page 310 |
| Performance considerations for your search queries | “Optimizing Search Queries” on page 312 |



| For information on: | Refer to: |
|---|---|
| Using the Audit Logging feature to audit library system usage | “Audit Logging” on page 323 |
| Using the Secure Document Delete feature to ensure that deleted documents cannot be accessed with disk recovery tools | “Secure Document Delete” on page 331 |
| Deleting outdated or inactive sessions | “Deleting Sessions” on page 334 |
| Performing regular backups of your library system's database and repositories | “Backing Up the Library System” on page 336 |
| Restoring the database that serves as a library system's property repository | “Restoring the Database” on page 339 |
| Restoring library system software | “Restoring Library Services” on page 340 |
| Restoring storage repositories | “Restoring Storage Repositories” on page 343 |
| Restoring archive repositories | “Restoring Archive Repositories” on page 345 |
| Restoring content search repositories | “Restoring Content Search Repositories” on page 346 |

Configuration And Maintenance

The property repository and content search repositories, like the storage repositories that are linked to it, have the potential of using up the storage space that you assign to it. For this reason, the library system is designed to assist you in monitoring and managing property repository and content search repository storage space. You should pay special attention to system-generated warnings concerning disk space usage. In response to these messages, you can use your database utilities to avert and, in some cases, recover from disk space problems that could otherwise cause suspension of normal operations for your users.

Any configuration or maintenance needs related to Verity K2 Admin will be done when reconfiguring library systems (such as moving library services from one server to another).

Storage Space Overflow

The property repository and content search repositories can overflow their assigned disk space in the following ways:

- (MS-SQL databases) Transaction data fills up the disk space assigned to the transaction log.
- Object property data fills up the disk space assigned to the property repository.
- A content search repository can reach its maximum allotted size.

The library system is set up to warn you of impending disk-space problems before they become critical enough to affect user operations. The property repository and (for Microsoft SQL Server databases) the associated transaction log is assigned two thresholds, one to warn you of pending disk-space overflow (initially set at 80% full) and the other to indicate that a critical “full” condition has been exceeded (initially set at 90% full). These thresholds are monitored by the reclaim daemon that runs automatically on each library system server. This daemon normally runs every five minutes as part of the version reclaim process. In addition, it checks the property server to see how much of the allotted disk space has been used by the property repository and (for Microsoft SQL Server databases) its transaction log.

Refer to [“Avoiding Property Repository Overflow” on page 225](#) for details on how to set the thresholds for the property repository and (for Microsoft SQL Server databases) the transaction log.

When the reclaim daemon finds that the property repository or (for Microsoft SQL Server databases) transaction log has reached the first threshold, the library system generates a warning message that is displayed on each library system server monitor. This message will be updated and displayed every time the reclaim daemon runs until more disk space is made available. A similar message is sent once to the message queue in the Administrators Group object and is also displayed on the administrative console each time an administrator opens a CS Explorer session.

If you do not respond to these threshold warning messages, the property repository or (for Microsoft SQL Server databases) transaction log may fill up until it reaches its “full” threshold, at which time the library system generates another message. Like the warning messages, the full threshold message is displayed on the monitor of each library system server that detects the condition, is sent to the Administrators Group object, and is also displayed on the administrative console each time an administrator opens a CS Explorer session. This time, however, each library system service is also *internally* suspended. When the services on a library system server are internally suspended, *Running* will still be the value of the Status property in the Storage Manager object, but all file-related operations, except those that you will need to clear more disk space, will be disallowed (for example, versions cannot be checked in or checked out, but you can back up the library system). Users will not be able to do any further work in

the library system until you free additional disk space for the property repository (and if necessary, for Microsoft SQL Server transaction log).

Some library system operations that do not access the Storage Manager (such as adding users or groups) may fill up the allotted disk space for the property repository beyond this “full” threshold before the appropriate message is actually displayed.

Avoiding Property Repository Overflow

To avoid error conditions caused by a full property repository or (for Microsoft SQL Server databases) transaction log, you can do the following:

1. Use CS Explorer regularly to check the System object and determine how much of the allotted disk space has been used by the property repository and (for Microsoft SQL Server databases) the transaction log.
2. Perform backups on a regularly scheduled basis.

(For Microsoft SQL Server databases) You can usually avoid filling up the transaction log by regularly running backups of the library system database engine. Each backup operation empties the transaction log and moves its contents into a backup file.

3. Modify the threshold settings in the System object.

Depending on your site's operational procedures, the default settings of 80% and 90% may not ensure adequate notification. For example, if procuring a larger server takes a long time at your site, you may want to lower the initial warning threshold so that you have plenty of advance notice that the property repository is filling up.

(For Microsoft SQL Server databases) You may want to lower the values of the Log Warning Threshold and Log Full Threshold. For example, if you have a small property repository but a large number of daily transactions, you may want to lower the thresholds for the transaction log so that you have plenty of time to respond and plenty of disk space remaining to perform the backup operations required to clear out the log.

If any of the threshold settings are exceeded, you (and all other administrators) will receive a warning message the next time you open a session.

If you type in a value of 0 for any one of these parameters, you will disable the associated threshold and no messages will appear for it. To avoid an unanticipated suspension of library system operations, you should use this 0 setting with great care.

4. Respond as soon as possible to threshold warning messages in one or more of the following ways:
 - (For Microsoft SQL Server databases) Perform a backup of the library system database engine to empty the transaction log (move its contents to a backup file).
 - Assign more storage space to the property repository.
 - (For Microsoft SQL Server databases) Assign more storage space to the transaction log.

If you do not have additional device space that you can assign on the property server, contact FileNet Support about other options, such as moving the database engine to a larger device or server.

5. To ensure that content search repositories do not reach their capacity, set the Default Indexed Version Limit in the System object to “No.”

Managing Storage Space

Because the library system enables you to store and access multiple versions of files over an extended period of time, a large amount of disk space may be required on the network servers. To ensure that the library system functions correctly, you should regularly monitor the available disk space. Since the library system is probably not the only application storing files on the network, you should devise methods for logging network traffic and storage requirements that are appropriate to your site. It makes no difference whether it is the library system itself or the other applications that use up the space on a particular disk. Once a storage server is full, it becomes impossible for the library system to store new version files on the storage repositories that are mapped to that disk, and library system users will encounter error conditions.

To avoid performance problems, you should not store more than 250,000 version files in a storage repository, even if the storage repository is not filled. Otherwise, the time to add or retrieve a version file will noticeably increase.

Storage Space and Storage Repository Overflow

A *storage repository overflow* occurs when no further version files can be added to the storage categories that are assigned to storage repositories on a particular library system server.

Storage repository overflow can occur in the following ways:

- A storage repository resides on a library system server that has become too full to accept a checked-in version file.
- An archive repository resides on a library system server that has become too full to accept an archived version file that is moving offline to make room for a newly checked-in version.
- A storage repository resides on a library system server that has become too full to accept a newly reclaimed version file from a reclaim repository.
- A storage repository is configured for Automatic Management and the maximum limit for the storage space or number of files for the storage repository, for which no Overflow Storage Repository is specified, is reached. Refer to “[Automatic Management of Storage Repositories](#)” on page 238 for details.

Such overflow conditions can occur either because no more disk space is available or because the storage repository (directory) has reached its maximum number of files.

In the first case, the library system will simply inform the user attempting to check in the version that the operation cannot be completed because there is insufficient disk space. Upon acknowledging this message, the user is returned to the display from which the checkin was begun. The Checkout object is unaffected and the file remains wherever it was located (on the user's local disk).

In the second case, the checkin completes and the version file is stored in the appropriate storage category on the assigned storage repository. However, the existing version file that was to be moved offline to accommodate the item's online limit remains on the storage repository because the server containing the archive repository is full. The user will receive a message indicating that there is insufficient disk space on the library system server containing the archive repository. If this problem is not addressed, the storage repository can also become so full that the disk space on its server is used up, causing another storage repository overflow condition.

In the third case, the library system will send the user a message indicating that the reclaim operation could not be completed because of insufficient disk space. The library system leaves the version file on the reclaim repository and periodically retries the reclaim operation until it can be completed.

In the fourth case, the library system will inform the user attempting to check in the version that the operation could not be completed because the maximum limit for the storage space or number of files for the storage repository is reached.

Avoiding Storage Repository Overflow

A *storage repository overflow* occurs when no further version files can be added to the storage categories that are assigned to storage repositories on a particular library system server.

To avoid an overflow of stored files in a storage repository or archive repository you can take the following steps:

1. Monitor storage space. Check the storage space used by the library system on a regular basis.

One way to monitor storage space is to set up a log and periodically check the storage space on each device that the library system accesses. You can run a check every day, week, or month, depending on the level of activity at your site.

If you choose to take this approach, you can get the server node and device names from the Storage Repository objects and you can find out how many bytes are free each week using the appropriate operating system command.

2. Archive regularly. Regularly remove (store, archive, or destroy) the versions in the archive repository.

To avoid archive repository overflow, you must have your own tools and methods for removing versions from the archive repositories. For example, you can store the repository files on tape, print them, or destroy them. It is also a good idea to monitor the storage space used by the archive repositories.

3. Adjust storage requirements. Adjust the configuration of the Storage Managers whenever necessary to better handle the version files created by your users.
4. Adjust the configuration of storage repositories for Automatic Management. Monitor the utilization of currently used storage repositories and always have an Overflow Storage Repository assigned to them. Refer to [“Automatic Management of Storage Repositories” on page 238](#) for details.

Preparing for a Possible Repository Failure

The library system has several built-in features that will help you recover from hardware and software failure. However, there are always cases where the loss of some element of the library system will adversely affect your stored data in ways that the library system cannot directly prevent or correct. For example, even though all version files are designated as Read Only in storage repositories and archive repositories, there is always the possibility that these files could be destroyed or corrupted beyond repair due to problems outside of the library system.

Preparing for a Storage Repository Failure

After a storage repository failure, you will usually have the following major concerns:

- Responding quickly so that the library system is not down for a long period of time.
- Replacing the damaged hardware and software.
- Recovering any lost version files.

To increase the speed with which you can recover from a storage repository failure and to take better advantage of the capabilities already built into the library system, we recommend that you take the following precautionary actions:

1. Always keep an appropriate replacement disk formatted and available.
2. Perform frequent library system backups of each storage repository. Obviously, the amount of work being done at your site will affect how often you must do backups. However, you should also consider how much data you can afford to lose and how long you can take to recover it in the event of a disk failure.
3. Create your own tool for restoring from your backup media any storage repository files that are missing after a storage repository is restored.

Preparing for an Archive Repository Failure

Although the version files lost from a failed or destroyed archive repository are typically those that are being removed permanently from online library system control, there may be cases where the versions need to be recovered (for example, because backup copies of the versions were never archived). To improve your ability to fully recover from an archive repository failure, we recommend that you take the following precautionary actions:

1. Keep a record of where the archived version files are located in your backup storage medium (for example, on which tape). Without such a record you will have to manually search through the archive medium for the missing files. It is up to you to recover the files themselves and ensure that your records are consistent with the library system.
2. Archive as quickly as possible those versions that have been moved to reclaimable archive repositories. Preventing a large number of version files from accumulating on the archive repositories will mean a quicker and easier recovery in the event of a disk crash.

Adjusting Storage Manager Configurations

As the library system and storage requirements expand, you may need to change the current repository configurations.

There are a number of methods (combinations of methods) you can use to do this, depending upon which is most appropriate to your site and circumstances:

1. Modify one or more storage categories to link to a storage repository on a server having more storage space available. This should spread the repository files more evenly across the network resources.
2. Create a new storage repository, assigning it to any available storage server, and modify as many storage categories as necessary to link to this new storage repository.
3. Add a new or larger device to a library system server, create a new storage repository that is mapped to the new device, and modify one or more storage categories to link to this new storage repository.

Make sure that you map new storage repositories to a disk that can be used by the library system indefinitely.

4. Install a larger storage device to replace the one on which the storage repository is currently located. You must manually back up the repository files on the existing device to a safe storage location on another device while the new, larger device is being installed. Then, restore the storage repository files to the newly installed device (making sure to maintain the same directory tree on the same server for each repository file).
5. Change the Default Storage Category property in the System object. This will not affect where the versions for existing documents are being stored, but it will affect where the versions of new documents are stored, assuming that most users cannot override this default value.

Configuring Databases

Note See “[Optimizing Oracle](#)” on page 108 for the database parameter settings required for proper library system operation, as well as general database performance guidelines.

Library system performance, affecting both servers and clients, is dependent on timely access to the property repository. The factors determining this access can change when one or more of the following conditions exists:

- User activity on the library system has significantly increased.
- The availability of database memory space allocated for the property repository has increased or decreased.
- The library system is incorrectly configured (for example, during the library system installation).

Reconfiguring database parameters requires shutting down and restarting both the database and the library system services in order for the changed settings to be used.

Refer to the following topics for guidelines about administering your library system database. Before you reconfigure database parameters, also refer to the documentation for your database.

- “[Setting the Database Connections](#)” on page 229
- “[Setting Memory Allocation](#)” on page 230
- “[Setting the Redo Logs Buffer Allocation](#)” on page 231
- “[Setting Database Buffers Allocation](#)” on page 232

Setting the Database Connections

Both library system clients and library system services require database-allocated connections (MS-SQL Server) or processes (Oracle). As the level of library system activity increases (generally as new groups and users are added to the library system), you may find that you need to increase the maximum number of concurrent, active client connections supported for the database engine.

For UNIX systems, increasing the number of connections/processes may also require increasing settings for kernel parameters.

You will need to ensure that enough connections remain for library system services (such as a Storage Manager), as well as the possibility of multiple library system sessions for a single client. In addition, Microsoft SQL Server customer-supplied databases may have other applications requiring database connections, and you will need to account for these also.

You control number of connections/processes with the following database-specific parameter:

| Database Platform | Parameter Name |
|----------------------|------------------|
| Microsoft SQL Server | user connections |
| Oracle | processes |

If system resources are at a minimum and the default number of client connections is more than adequate for the library system, you could decrease this number. An example of where this may be appropriate is when the total number of Storage Manager processes/threads available is less than the number of client connections allowed, and you do not have the servers available to add more Storage Managers.

To set the maximum database connections supported for Microsoft SQL Server databases:

1. Start Microsoft SQL Server Enterprise Manager.
2. Locate the *user connections* parameter and increase this as necessary.
3. Increase the memory space available. Refer to “[Setting Memory Allocation](#)” on page 230 for more information.
4. Shut down all library system services.
5. Shut down and restart the database.
6. Restart all library system services.

To set the maximum processes supported for the library system database instance on Oracle databases:

Note Before changing the value of the processes parameter, refer to “[Optimizing Oracle](#)” on page 108 for specific setting requirements.

1. Using a text editor, open your *ORACLE_SID.ora* file.
2. Locate the processes parameter and increase this based on the values recommended in the commented text.
3. Increase the memory space available, if necessary. Refer to “[Setting Memory Allocation](#)” on page 230 for more information.
4. Shut down all library system services.
5. Shut down and restart the library system database.
6. Restart all library system services.

Setting Memory Allocation

Each time a library system client or service accesses the property repository, a certain amount of memory is required on the property server to handle the client connection or service request. For this reason, whenever you increase the setting for the number of database connections supported, you should also verify that enough memory is available to the property repository.

Note For information about RAM requirements, refer to [CS 5.5.0 Windows Oracle Platform Requirements at FileNet Worldwide Customer Support](#).

Note By default, Microsoft SQL Server dynamically controls its own memory usage and will automatically acquire the memory it needs as client connection or service requests are received. In this configuration, as the number of client connections and service requests increase, so will SQL Server's memory usage, which can be seen to consume a large amount of physical memory and then fluctuate. This memory usage is normal and generally, the default configuration works well for library systems. However, if you have additional applications running on your server, which have heavy memory requirements (library system services do not), or if the MSSearch service is used

(Content Services does not use it) you may need to adjust SQL Server's maximum memory allowance, either out of necessity or to improve performance by avoiding excessive page faults. Refer to “Server Memory Options” and “Memory Architecture” in *Microsoft SQL Server Books on Line* for more information.

The way memory allocation is controlled is database-specific.

| Database Platform | Memory Control |
|----------------------|--|
| Microsoft SQL Server | Memory tab of SQL Server Properties Dialog Box |
| Oracle | shared_pool_size parameter |

Note Before changing the value of the shared_pool_size parameter, refer to “Optimizing Oracle” on page 108 for specific setting requirements.

To set Microsoft SQL Server's memory configuration:

1. Shut down all library system services. Refer to “Controlling Library System Services” on page 198 for information about stopping, starting, and checking the status of services.
2. Start Enterprise Manager.
3. Expand the Microsoft SQL Servers tree and the SQL Server Group.
4. Under SQL Server Group, select your property server, right-mouse click and select Properties.
5. On the Memory tab of the SQL Server Properties dialog box, adjust your server's memory allowance as necessary. Refer to *Microsoft SQL Server Books on Line* for information on using these features.
6. Shut down and restart the database.
7. Restart all library system services.

To set the shared memory space supported for Oracle databases:

1. Using a text editor, open your **SPFILEDATABASENAME.ORA** (Oracle9i or Oracle 10g) file.
2. Locate the **shared_pool_size** parameter and increase this based on the values recommended in the commented text.
3. Shut down all library system services.
4. Shut down and restart the library system database.
5. Restart all library system services.

Setting the Redo Logs Buffer Allocation

All library system database transactions are written to logs. For Oracle databases these are online redo logs. As the level of activity on the library system expands, the number of transactions may increase beyond the capacity of the initial logging configuration setting. This is controlled by the log_buffer memory buffer parameter.

To set the redo log buffers allocation:

Note Before changing the value of the `log_buffer` parameter, refer to “[Optimizing Oracle](#)” on page 108 for specific setting requirements.

1. Using a text editor, open your `SPIFILEDATABASENAME.ORA` (Oracle9i or Oracle 10g) file.
2. Locate the `log_buffer` parameter and increase this based on the values recommended in the commented text.
3. Shut down all library system services.
4. Shut down and restart the library system database.
5. Restart all library system services.

Setting Database Buffers Allocation

Library system database transactions are cached in a buffer (using 4 Kb block sizes) in order to maximize database I/O. If the buffer size is insufficient, as the level of activity on the library system expands, the number of transactions may impact performance. Buffer size is determined by the setting for the `db_block_buffers` memory buffer parameter.

To set the database buffers allocation:

Note Before changing the value of the `db_block_buffers` parameter, refer to “[Optimizing Oracle](#)” on page 108 for specific setting requirements.

1. Using a text editor, open your `SPIFILEDATABASENAME.ORA` (Oracle9i or Oracle 10g) file.
2. Locate the `db_block_buffers` parameter and increase this based on the values recommended in the commented text.
3. Shut down all library system services.
4. Shut down and restart the library system database.
5. Restart all library system services.

Storage Repositories


Storage repositories are the directories or partitions you create for storing documents added to a library system.

When a storage repository is added to a library system, a corresponding Storage Repository object instance is created. It contains several descriptive properties, including the system-assigned storage repository ID, the directory path, and the network name of the server on which the storage repository resides.

The Storage Repository object also contains a list of the storage categories that are assigned to the storage repository. Each document version is checked in to the library system according to a storage category that is assigned to the associated Item object. This storage category is mapped to a storage repository directory that will receive and store the file. A single storage repository may contain any number of storage categories, and each storage category may contain the versions of any number of documents.

In addition to a Storage Category property, each Item object contains an Online Limit property that controls how many versions of a document can be stored on the storage repository. When this online version limit is reached, a subsequent checkin will prompt the library system to remove an appropriate older version from the storage repository to make room for the newer version being created. The version that is moved out of the storage repository is stored in an archive repository, where it will remain until you remove it for final archiving or destruction outside of the library system.

The following illustration shows the relationship of the storage repositories to the Storage Manager and the storage categories and how these elements work together to control and track your versions.



Storage Manager, Repositories, and Categories

Because the value of the Storage Category property is modifiable in the Item object, it is conceivable that different versions of a document could be stored under different storage categories and, in some cases, on more than one storage repository. The Online Limit value will always control the total number of versions of the document that are stored online, regardless of the storage category assigned.

During the process of installing the initial components for a particular library system, Storage Repository 1 (**shelf001**) is created automatically. An initial storage category (the default is Documents) is also created and assigned to Storage Repository 1. You will probably add other storage repositories during the initial configuration process. You can also add them later to meet growing storage demands and to designate storage areas on additional library systems.

After a storage repository has been added, you can modify its object directly only by entering new descriptive data to more clearly identify its size or purpose. You cannot delete the storage repository or its object, although you can assign some or all of its storage categories to other storage repositories.

You can specify the maximum size for a storage repository or the maximum number of files that a storage repository can store. You can also configure a storage repository to automatically switch to a predefined overflow storage repository when limits specified for the storage repository in terms of maximum size or maximum number of files is reached.

Adding a Storage Repository

To add a storage repository:

Note For tool-specific instructions on how to add a storage repository, refer to the *CS Explorer Online Help*.

1. Determine the network name of the library system server and device to which you want to add the storage repository.
2. (UNIX servers only) Recommended: Create a dedicated partition for each storage repository you want to add to ensure the availability of adequate disk space for the version files that will be stored.
3. Navigate to the storage repositories for this library system.
4. Add a storage repository.
5. Enter property values for the new storage repository. Required properties include:
 - Device
 - Server
6. Save your changes.
7. If a message indicates that disk space is not available for the new storage repository, you must do one or more of the following:
 - Delete (or store offline) some of the files or directories from the current disk.
 - Obtain a larger disk.
 - Add the storage repository to a different existing library system server (specify a different value for the Server and/or Device property).

- Set up an additional Storage Manager on another server.

Modifying a Storage Repository

To modify a storage repository:

Note For tool-specific instructions on how to modify a storage repository, refer to the *CS Explorer Online Help*.

1. Navigate to the storage repositories for this library system.
2. Select the storage repository you want to change.
3. Modify the property values you want to change. Modifiable properties include:
 - Comment
 - Device (Refer to “[Moving a Storage Repository](#)” on page 235 for instructions on how to modify this property.)
 - Network Path
 - Server
 - Max Disk Space (GB)
 - Max Space Warning Threshold (%)
 - Max Files
 - Max Files Warning Threshold (%)
 - Overflow Storage Repository

Note Changes to the attributes for Automatic Management of Storage Repositories are effective only when the Storage Manager is restarted. Refer to “[Automatic Management of Storage Repositories](#)” on page 238 for guidelines on using these properties.

You can also access storage categories from the Storage Repository object. Refer to “[Modifying a Storage Category](#)” on page 241 for details.

4. Save your changes.

Moving a Storage Repository

To move a storage repository:

Note For tool-specific instructions on how to move a storage repository, refer to the *CS Explorer Online Help*.

1. Put the library system in single-user mode. A library system is in single-user mode when:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.

- On the target device, recreate the directory path that contains the storage repository files. The storage repository files can be found in:

| OS Platform | Location |
|-------------|---|
| Windows | <i>Drive:\filenet\LibrarySystem\Stacks\Shelfxxx</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/stacks/shelfxxx</i> |

- Move the storage repository files from the original device to the target device.

Note You must delete the storage repository files from the original device. Failure to do so may cause unexpected results.

- Modify the Device property for the storage repository to point to the new target device.
- Enable logons by setting the Logons Enabled? property in the System object to Yes.

Moving Versions Between Storage Repositories

At some point you may want to move document versions between storage repositories. You can do this with the DMU utility (see the DMU description in the Utilities section for complete details), using a transfer repository and a transfer category. A *transfer repository* is an archive repository whose only purpose is to serve as intermediate storage between the source and target (destination) storage repositories. A *transfer category* is an archive category assigned to the transfer repository. Neither the transfer repository nor the archive category should be used for any purpose other than moving versions between storage repositories.

To move versions between storage repositories:

Note Perform the version move when there is little or no user activity on the library system. Otherwise, users' attempts to view or check out documents could interfere with the version movement.

- Create a new archive repository to serve as the transfer repository. Set its Reclaimable? property value to Yes and its Delete Upon Archive? property value to No. You will need to know this ID in [step 6](#) and [step 8](#) of this procedure. To help you find it, run the DMU utility, which lists the IDs of all the archive repositories in the library system:

```
DMU -A propertyserver librarysystem logfile
```

Note For tool-specific instructions on how to add an archive repository, refer to the CS Explorer Online Help.

- Create a new archive category, called transfer category in this procedure, and assign it to the transfer repository created in [step 1](#).

Note For tool-specific instructions on how to add an archive category, refer to the CS Explorer Online Help.

- If any document whose versions you wish to transfer is in a replicated folder, move that document out of the replicated folder. Be sure to allow enough time for Replication Services to process the dereplication before continuing at [step 4](#). (You will return these documents to their replicated folders in [step 8](#).)
- Run the ARCHDATE utility to archive each version that you wish to transfer. (See the ARCHDATE description in the Utilities section for details.) ARCHDATE can archive versions by date or by stored search. You must use the ForceArchCat option to override each version's default archive repository. You must also set

Substitute ArchCat to the transfer category you have created. You can run ARCHDATE multiple times with different searches.

Note You must not use the archiving features of Content Services client software (such as IDM Desktop) in place of ARCHDATE. If you do, the versions will be sent to their default archive repositories, not the transfer repository.

5. Run the DMU utility to see the IDs of the storage repositories in the library system that contains the versions you want to move. Choose the ID of the storage repository you want as the target. You will need to know this ID in [step 6](#) and [step 8](#).

```
DMU -S propertyserver librarysystem logfile
```

6. Run the DMU utility in test mode to identify all the versions that you have archived to the transfer repository in [step 4](#). In test mode DMU does not actually move the versions; it just shows you their item and version IDs:

```
DMU -t propertyserver librarysystem logfile transfer_repository_ID  
target_repository_ID
```

7. If all the versions you wish to move to the target repository are listed in the output of the DMU utility, then continue at [step 8](#); otherwise return to [step 3](#).
8. Run the DMU utility again, this time in execute mode (you must have Admin access to the library system), to actually move the versions from the transfer repository to the target repository:

```
DMU -x propertyserver librarysystem logfile transfer_repository_ID  
target_repository_ID admin_user admin_pwd
```

As the DMU utility runs in execute mode, it updates the storage repository of each version in the transfer repository to be the target repository. It then reclaims each version and moves its associated file into the target repository. After each reclaim, DMU looks up the archive repository currently associated with the archive category of the item to which the version belongs and resets the archive repository of the version to match.

Note If the library's archive category/repository associations have changed over the lifetime of the library system, the archive repository after the reset may not be the same as before the version move.

9. If you moved any documents out of replicated folders in [step 3](#), move them back into their folders.
10. For each version that DMU reclaims in [step 8](#), it leaves a copy of the file in the transfer repository. Since these files are no longer needed, you should remove them.

Setting Up a Direct-Document-Access Repository

Library systems support the option to set up direct document access (DDA) on storage repositories, which allow some user interfaces, such as IDM Desktop, to have faster, more direct viewer access to version files stored in those storage repositories. To set up DDA, you must share the path and give read access to a storage repository on the server where the documents are located. You must also enable DDA at the user interface, such as IDM Desktop or the CS Web Server (see the *IDM Desktop Administrator's Guide* or the *Web Services Administrator's Guide* for details).

Be aware that making files on storage repositories available for direct viewing from client programs raises access security issues that may be important at your site. If you would prefer to maintain library-system file-access security checks (and can accept slightly slower viewer access), then do not set up DDA.

To share the path and give read access to a storage repository:

1. From Windows Explorer, right-click the shelf directory (for example, d:\filenet\library\stacks\shelf001) containing the version files for which you want to set up DDA. Then choose Properties from the pop-up menu. We recommend that you share at a level no higher in the directory path than is necessary.
2. In the Properties dialog box, select the Sharing tab and click the Shared As option button. Type a share name (the default is the name of the directory) in the Share Name box. Then click Permissions.
3. In the Access Through Share Permissions dialog box, click Read in the Type of Access list, click OK, and then click OK in the Properties dialog box. This gives read access to the files in the shelf directory.
4. Start CS Explorer and log in as an administrator to the library where you want to set up DDA. In the Tree view of the library, double-click the Storage Repositories icon and right-click the storage- icon corresponding to the shelf folder you selected in [step 1](#), and select Properties. For example, if you selected shelf001, then right-click the storage repository named *I*.
5. In the Storage Repositories property sheet, type a Network Path value corresponding to the directory location of the shelf folder. For example, if the storage repository resides on a machine named server and you specified a share name shelf001 in [step 2](#), then type \\server\ifilenet\library\stacks\shelf001. Then click OK.

Automatic Management of Storage Repositories

Automatic management of storage repositories brings automation to CS administration, leading to more effective disk space management and better performance. Using automatic management of storage repositories, administrators can configure how much space a particular storage repository (in other words directories or partitions) can utilize or how many documents can be located in a storage repository. Once the maximum limit for the storage space or number of files is reached, the system switches to the specified overflow storage repository and also sends out a notification to the administrators. Overflow storage repositories can also be configured to have their own overflow storage repositories, thus allowing multiple levels of redirection.

Storage repository for which one of the thresholds has been reached and that has switched to an overflow repository is marked as *frozen* and is unavailable for future use. However, the *last* storage repository (storage repository for which the overflow repository is not specified) is not marked as frozen when it becomes full. Instead, an error message is returned when an attempt is made to add a new file to this storage repository.

The following actions cannot be performed on a storage repository that has been marked as frozen:

- Assigning the storage repository as an overflow storage repository for another storage repository.
- Assigning the storage repository as the default storage repository for a storage category.
- Changing the overflow attributes for the storage repository.

The CS System sends out notifications to the administrators under the following conditions:

- The maximum limit for the storage space or number of files for a storage repository is reached and the system has switched to the specified overflow storage repository.
- Overflow storage repository has not been specified for a storage repository and utilizations have exceeded the specified warning threshold.

Note Warning message is sent once in a day or whenever the storage manager for the repository is started.

The following attributes need to be specified for configuring automatic management of a storage repository:

- **Max Disk Space (GB)** - Indicates the maximum disk space (in Gigabytes) that the storage repository can occupy. Once this limit is reached, the system switches to the Overflow Storage Repository, if specified. This value should be greater than 0 and less than 2048. Once specified, this value cannot be decreased.
- **Max Space Warning Threshold (%)** - Indicates the percentage of **Max Disk Space** that must be used before a warning threshold message for disk space usage is issued to the Administrators group. If this value is not specified, it will be defaulted to 80.
- **Max Files** - Indicates the maximum number of documents that can be stored in the storage repository. Once this limit is reached, the system switches to the Overflow Storage Repository, if specified. Once specified, this value cannot be decreased.
- **Max Files Warning Threshold (%)** - Indicates the percentage of **Max Files** that must be reached before a warning threshold message for file count is issued to the Administrators group. If this value is not specified, it will be defaulted to 80.
- **Overflow Storage Repository** - ID of the next storage repository to be used when the limits specified for this storage repository are reached. The overflow storage repository assigned should not cause circularity in the order of overflow repositories and should not be frozen.

WARNING If a storage manager is abruptly shutdown, it takes a little extra time than usual to start up as it calculates the current disk space usage and file count of each storage repository that has been enabled for automatic management.

Storage Categories

Storage categories are groupings of documents that have common storage requirements. You define the storage categories that will best suit a particular library system. For example, if you are administering a library system largely oriented to word-processing, you could define its storage categories as Documents, Graphics, and Memos. On the other hand, if you are administering a library system oriented to financial data, you could define its storage categories as Spreadsheets, Charts, Schedules, and Miscellaneous.

Your job is to link the storage categories to the appropriate storage space, as represented by the storage repositories you assign. For example, if you were administering the word-processing library system mentioned above, you could assign the Documents and Graphics storage categories to a storage repository on a device with a large amount of space, and assign the Memos storage category to a storage repository on a device with less space allocated.

You'll need to estimate the size of the storage repositories associated with each storage category. One approach is to see how much space, collectively, is presently used for each document category. Your space estimate must accommodate the potential growth of each storage category. The best approach to this problem is to study the past storage requirements of the users. While doing this, you may find that you want to change the names you decided on for the storage categories or add some different ones.

Note A storage category cannot be renamed.

If your library system has a Content Search Manager, you will also assign a content search repository to some storage categories. Documents in those storage categories can be indexed for content searching based on the Assigned Content Search Repository property value in the Storage Category object. By default, new storage categories are not assigned to a content search repository.

To reduce network traffic during indexing processes, we suggest that, whenever possible, content search repositories be assigned to storage categories associated with storage repositories existing on the same host as the Content Search Manager.

When a document is added to a library system, a default storage category is assigned to it. However, an administrator can change the assigned storage repository for a storage category at any time. New versions of the documents using that storage category are then stored in the newly assigned storage repository.

You can add up to 999 storage repositories and assign as many storage categories as appropriate for the library system, thereby giving you full control over how you distribute the document storage load across available hardware resources and individual library systems.

Each storage category has its own object instance with Storage Category (name), Comment, and Assigned Storage Repository properties. You provide values for these properties when you add the storage category to a library system. Most sites begin by adding an initial set of storage categories and then add more as needed. Although you can add new storage categories at any time and have as many as you want, you cannot remove or rename them. This restriction is enforced because an existing storage category may have been used in many places throughout the library system.

When users add a document to the library system, they simply select a category that best describes that document. If you find over time that particular storage categories require more or less storage space, you can reassign them to different storage repositories, adjusting the network load as you see fit.

Adding a Storage Category

To add a storage category:

Note For tool-specific instructions on how to add a storage category, refer to the *CS Explorer Online Help*.

1. Navigate to the storage categories for this library system.
2. Add a storage category.
3. Enter property values for the new storage category. Required properties include:
 - Storage Category (name)
Double-check the spelling of the storage category name. Once the storage category has been added to the system, you cannot modify the name.
 - Assigned Storage Repository
If you do not specify a content search repository to this storage category, the library system will use content search repository 1 when indexing the item.
4. Save your changes.

Note The assigned storage repository should not be frozen. **Storage Repository Full?** property for the storage repository should not be set to **Yes**.

Modifying a Storage Category

To modify storage categories:

Note For tool-specific instructions on how to modify a storage category, refer to the *CS Explorer Online Help*.

1. Navigate to the storage categories for this library system.
2. Select the storage category you want to change.
3. Modify the property values you want to change. Modifiable properties include:
 - Assigned Content Search Repository
 - Assigned Storage Repository
 - Comment
4. Save your changes.

Archive Repositories

Since each document in a library system has one or more versions, the amount of data in a system's storage repositories could eventually grow beyond reasonable limits. Each library system uses parameters you establish to begin the process of moving offline some (or all) document versions.

Archive repositories are the directories or partitions you create to temporarily store the versions that are moved offline by the archiving process. You can manually archive versions, or you can specify an online version limit and let the library system handle the process for you automatically. In most cases, the result is that the versions are transferred to archive repositories, directories that you create to temporarily hold these files before you delete or move them to offline storage.

Files that are to be archived or placed in offline storage are reclaimable, because they may be retrieved and returned to online library system control. These files will be moved offline through archive categories that are assigned to reclaimable archive repositories. Other files that will eventually be destroyed are nonreclaimable, because they cannot be easily retrieved (if at all) for return to online library system control. They are moved offline through archive categories that are assigned to nonreclaimable archive repositories. You can set up nonreclaimable archive repositories to immediately delete, rather than temporarily store, versions archived to it; but to do so, you must set the Delete Upon Archive? property value to Yes when you first add the archive repository to the system.

Reclaimable archived versions, if not destroyed, can be moved back to the library system and accessed by authorized users. To facilitate this reclaim process, the library system creates the following additional directories when you add an archive repository to a library system:

- The reclaim command repository contains text command files that are created to identify versions whenever users submit version reclaim requests.
- The reclaim repository is where you place copies of the offline versions that you have located in your backup media.

Refer to [“Reclaiming Offline Versions” on page 244](#) for further details on how versions are reclaimed.

When an archive repository is added to a library system, a corresponding Archive Repository object is created in the associated library system. It contains several descriptive properties, including the system-assigned archive repository ID, the directory path, and the network name of the server on which the archive repository resides.

The Archive Repository object also contains a list of the archive categories that are assigned to the archive repository. Each version that is taken offline (transferred from the storage repository) is moved according to an archive category. This category determines which archive repository (**dflexxx directory**) will receive and temporarily store the file as it awaits final offline archival or (manual) deletion. Any number of archive categories may be linked to a single archive repository.

Files that are to be archived or placed in offline storage are reclaimable, because they may be retrieved and returned to online library system control. These files will be moved offline through archive categories that are assigned to reclaimable archive repositories. Other files that will eventually be destroyed are nonreclaimable, because they cannot be easily retrieved (if at all) for return to online library system control. They are moved offline through archive categories that are assigned to nonreclaimable archive repositories. You can set up nonreclaimable archive repositories to immediately delete, rather than temporarily store, versions archived to it, but to do so you must set the Delete Upon Archive? property value to Yes when you first add the archive repository to the system.

Archive Repository 1 is created automatically during the process of installing the initial services for a particular library system. This archive repository is intended for reclaimable versions. An initial archive category (Trash) is also created and assigned to it. You will probably add other archive repositories during the initial configuration process. You can also add them later to meet changing file-storage demands and to designate storage areas on additional servers that may be installed at your site.

After an archive repository has been added, you cannot remove it. You can, however, modify its Comment property by entering new descriptive data to more clearly identify its size and purpose.

Moving Versions Offline

Versions can be moved offline onto archive repositories in two ways:

- **Automatic Archiving.** Occurs according to the online version limit set by the user for a particular document. When the checkin of a new version causes the online version limit to be exceeded, the “oldest” version (based on lowest version number) is moved offline. This automatic archiving can affect a version only if the Permanent Version? property in its Version object has a value of No. For any particular document, the latest version is also exempt from automatic archiving.

Note Although versions marked Permanent cannot be moved offline during automatic archiving, they may be counted against the online version limit if the value of the Permanent Versions In Limit? property in the System object is set to Yes.

- **Manual Archiving.** Occurs when any user with at least Owner access rights to the Item object chooses the Archive command to request that a version be moved offline. The value of the Permanent Version? property in the Version object does not affect manual archiving.

Note FileNet now provides a separately priced Archive to Image Services product as an add-on to CS. The product allows you to archive and reclaim CS-stored content into and out of an existing IS library. You can configure this functionality to work in conjunction with the standard CS auto-archiving feature. Contact your FileNet sales representative for details. See the product documentation at [FileNet Worldwide Customer Support](#).

The archiving process affects only the versions themselves, not the associated object instances. All Version objects remain online at all times, regardless of the status of the version they describe.

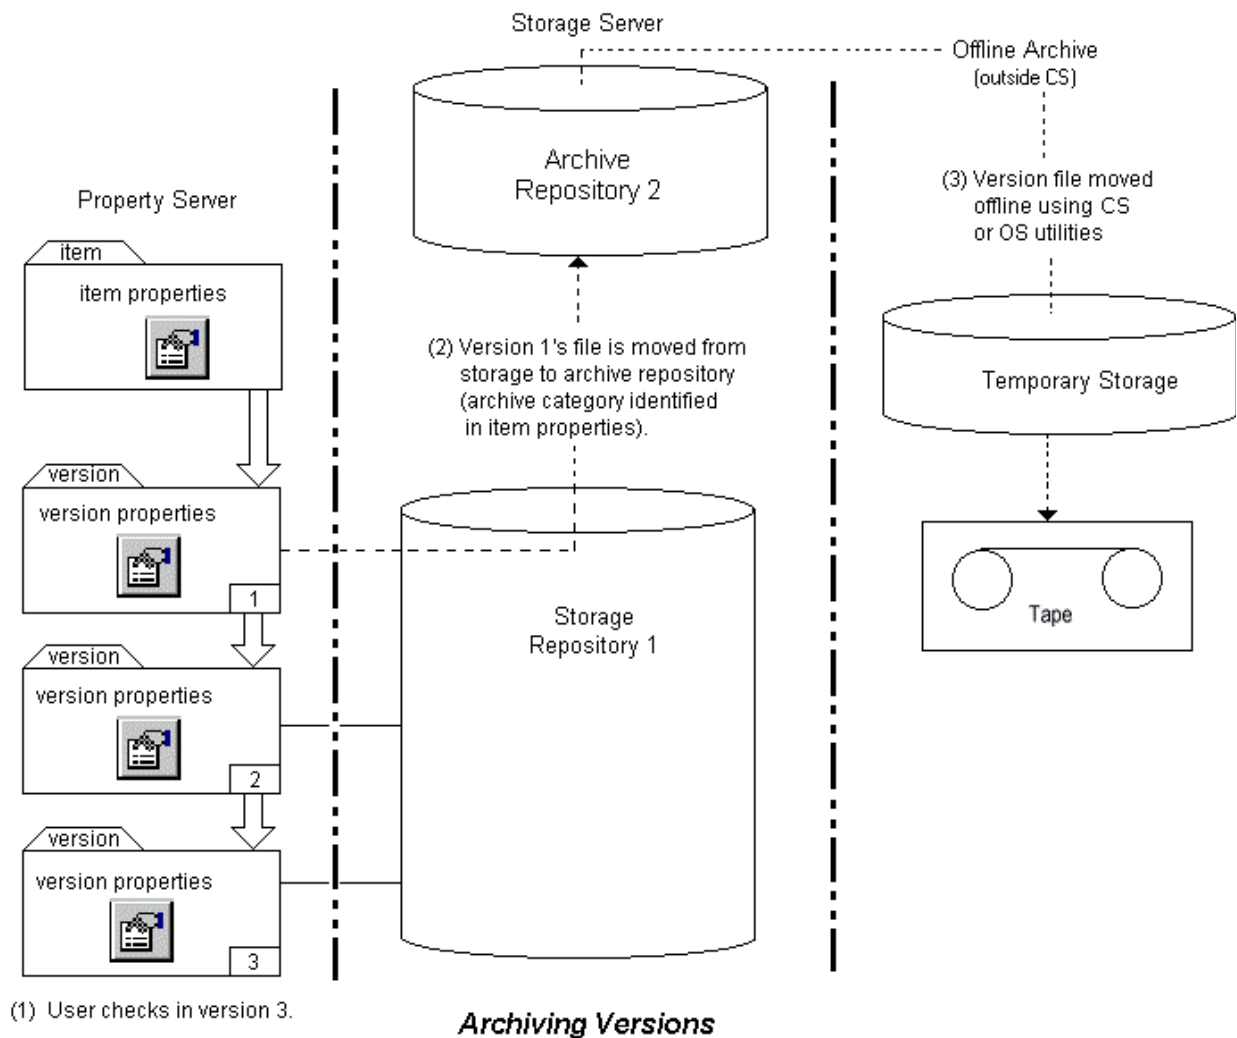
The following sequence of events takes place during version archiving:

1. Whenever an automatic or manual archiving request is issued, the library system removes the appropriate version from the storage repository, leaving only the Version object in the property repository as a historical record.
2. The library system uses the Storage Repository ID and Encrypted File Name properties in the Version object to identify the storage repository number and the system-generated file name for the version.
3. The library system identifies the correct archive category for the version (for example, Archive To Magnetic Tape) by reading the value from the Item object. Each archive category is assigned to a particular archive repository.
4. Using the information obtained in [step 2](#) and [step 3](#), the library system moves the version from the storage repository to the correct archive repository and changes the value of the version's Availability Status property to Offline (Reclaimable) or Offline (Nonreclaimable), depending on the type of archive repository used. Also, the library system updates the Archive Date property in the Version object to reflect the date of the archiving

and the archive repository ID to reflect the archive repository used. Although the version is now offline from the user's standpoint, it still must be stored or destroyed.

Since a library system moves the version only as far as the archive repository and does not interact physically with archiving or disposal media, you must use your normal network maintenance tools and procedures to move the file totally offline and out of the library system's control. You can specify the offline location by entering the tape/disk label or other identifying number as the value of the Offline Location property in the Version object.

The following figure illustrates the steps in the archiving process for moving versions offline automatically. The document in this example has an online version limit of 2, and versions 1 and 2 are flagged as not permanent.



Reclaiming Offline Versions

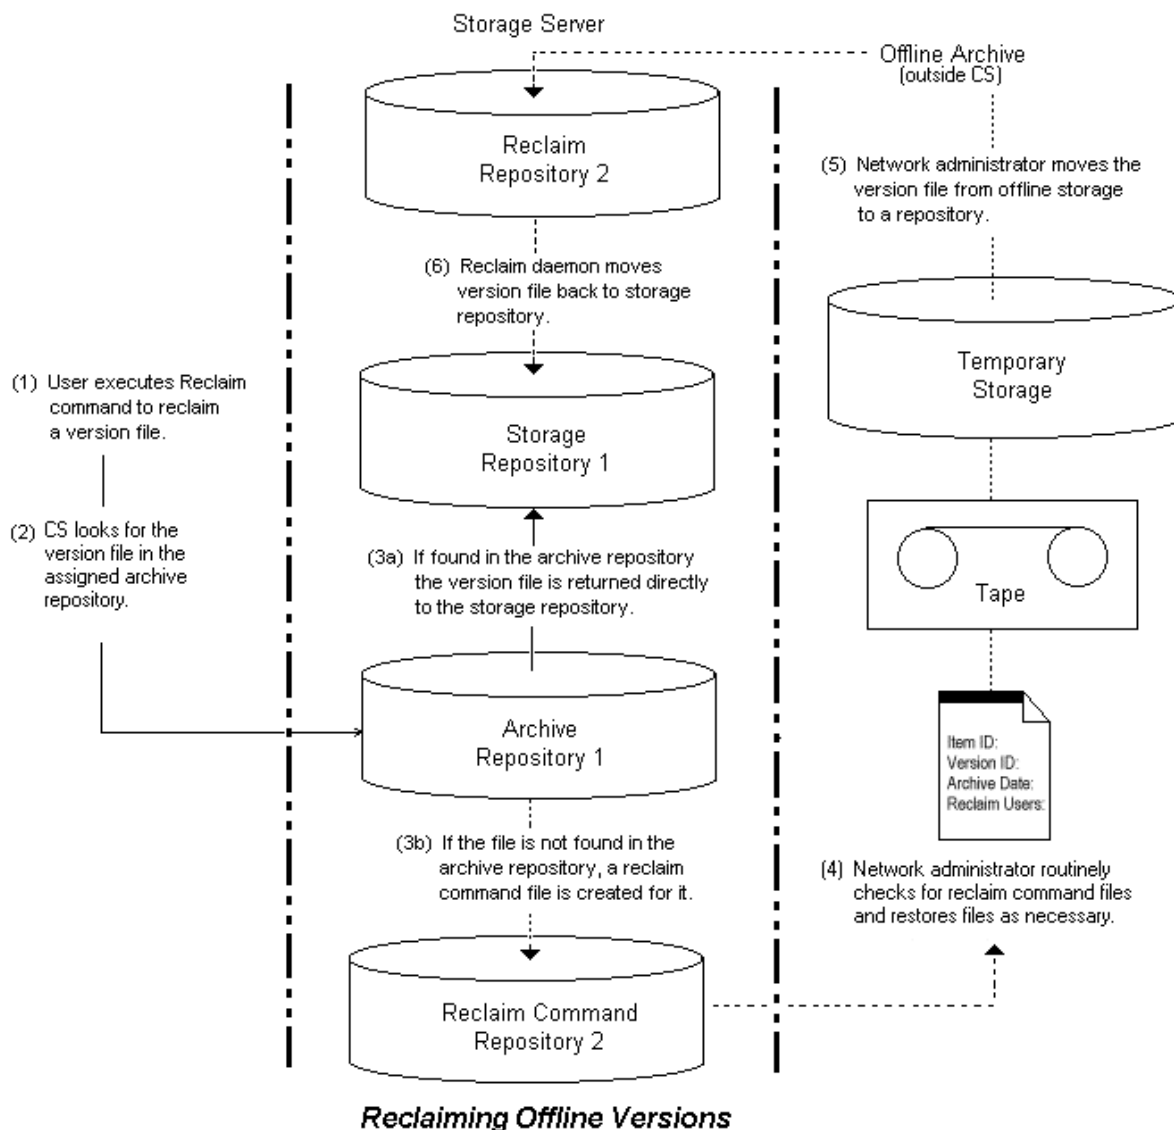
As mentioned earlier, each time you add an archive repository (**dflexxx**), the library system creates two corresponding directories for reclaiming versions for a particular library system:

- A reclaim command repository (**rcmdxxx**) for holding descriptions of the requested versions.

- A reclaim repository (**rfilexxx**) for temporarily holding versions immediately after they are reclaimed.

As an example of how a version is reclaimed using these various directories, let's assume that a version found on storage repository 1 (**shelf001**) was flagged for automatic archiving, moved onto archive repository 2 (**dfile002**), and then transferred out of the archive repository with your local tools onto offline storage medium (for example, magnetic tape).

The following figure illustrates the steps in the version reclaim process. Once this process begins, the version is either immediately reclaimed from the archive repository (step 2a) or the file must be retrieved from the offline storage medium (steps 2b through 5).



Processing a User's Reclaim Request

To reclaim an offline version, a user must have the necessary access rights to the version's parent item. The setting of the Min Reclaim Access Level property in the System object controls what access rights are required for a user

to reclaim a version. If the user's access rights are adequate, the reclaim request is accepted. You must then assist in processing it to ensure that the version is brought back online.

Note Items that are reclaimed are moved from an archive to the storage repository, they are not copied, so the item is deleted from the archive once it is successfully back in storage.

To process a user's reclaim request:

1. After a user has requested that a version file be reclaimed, the library system responds by checking the archive repository assigned to the version and performing one of the following steps:

- If you have not yet removed the version's file from its assigned archive repository (for example, **dfile002**), the library system moves it back to storage repository (for example, **shelf001**) as indicated by the item's storage category. The library system takes this step even if the archive repository is identified as nonreclaimable. It also changes the Availability Status property in the Version object from Offline to Online (indicating to the user that the version is once again available) and sets the Permanent Version? property to Yes.

This ends the reclaim process for any version file that has not been removed from its archive repository. The remaining steps apply only to a version file that has been removed from its respective archive repository.

- If you have removed the version's file from the archive repository, the library system places a command file in the associated reclaim command repository (for example, **rcmd002**) and changes the value of the Reclaim Pending? property in the Version object to Yes.

2. To process the reclaim requests for versions that have been removed from the archive repository, you must routinely check the reclaim command repositories (**rcmdxxx**) and read the text command files that may be stored there. The command file contains information such as the version ID, item ID, file name, archival date, date of the reclaim request, and name of the requesting user.

If you have already deleted the file because it was assigned to a nonreclaimable archive repository, you should inform the user who requested the reclaim.

3. Using your own network tools, retrieve the requested version file from offline storage and copy it from the backup medium onto the correct reclaim repository (in this example, **rfile002**). Retrieving the version may be fairly simple if a value has been entered in the Offline Location property in the Version object to identify the tape or disk on which the file is stored.

4. The library system reclaim daemon tool scans the reclaim repository looking for reclaimable files. Within minutes after the file is put onto the reclaim repository, the library system moves the file to the version's storage repository and deletes the command file from the reclaim command repository.

The library system also modifies the Version object, changing the value of the Availability Status property to Online, the value of the Reclaim Pending? property to No, and the value of the Permanent Version? property to Yes. In addition, a message indicating that the version file is now reclaimed and available is added to the Messages property in the User object of the person who requested the reclaim. The reclaimed version is marked as a permanent version, exempting it from future automatic archiving. If a reclaimed version causes the item's online version limit to be exceeded, the library system will attempt to move two versions offline during the next check in of a version.

Adding an Archive Repository

The following procedure shows you how to add an archive repository. This procedure is usually performed immediately following installation, but it can be used whenever archival and retrieval requirements change.

When adding a new archive repository, you can specify whether it will handle reclaimable or nonreclaimable files. You can use this feature to indicate whether versions moved offline through a particular archive repository can be easily reclaimed from their offline locations. Archive Repository 1, which is added during the installation of a library system, is considered a reclaimable archive repository.

After the archive repository is added to the system using the following procedure, the only change you can make to the object is to modify the text for the Comment property. To modify this property, use the procedure described in “[Modifying an Archive Repository](#)” on page 248.

To add an archive repository:

Note For tool-specific instructions on how to add an archive repository, refer to the *CS Explorer Online Help*.

1. Determine the name of the library system server and the name and storage capacities of available library system server directories to which you want to add archive repositories.
2. Navigate to the archive repositories for this library system.
3. Add an archive repository.
4. Enter property values for the new archive repository. Required properties include:
 - Delete Upon Archive?
 - Device
 - Reclaimable?
 - Server

Do not add an archive repository to removable media, such as a floppy disk.

5. Save your changes.

This step adds the archive repository to the library system and creates the corresponding reclaim command repository (**rcmdxxx**) and reclaim repository (**rfilexxx**) for use during reclaim operations.

6. If a message indicates that disk space is not available for the new archive repository, you must do one or more of the following:
 - Delete (or store offline) some of the files or directories from the current disk.
 - Obtain a larger disk.
 - Add the archive repository to a different existing library system server (specify a different value for the Server and/or Device property).
 - Set up an additional Storage Manager on another server.

Modifying an Archive Repository

To modify an archive repository:

Note For tool-specific instructions on how to modify an archive repository, refer to the *CS Explorer Online Help*.

1. Navigate to the archive repositories for this library system.
2. Select the archive repository you want to change.
3. Modify the property values you want to change. Modifiable properties include:
 - Comment
 - Device (Refer to “Moving an Archive Repository” on page 248 for instructions on how to modify this property.)
 - Network Path
 - Server

You can also access archive categories from the Archive Repository object. Refer to “Modifying Archive Categories” on page 251 for details.

4. Save your changes.

Moving an Archive Repository

Note For tool-specific instructions on how to move an archive repository, refer to the *CS Explorer Online Help*.

To move an archive repository:

1. Put the library system in single-user mode. A library system is in single-user mode when:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
2. On the target device, recreate the directory path that contains the archive repository files. The archive repository files are located in:

| OS Platform | Location |
|-------------|--|
| Windows | Drive:\filenet\LibrarySystem\Stacks\Dfilexxx |
| UNIX | \$IDMDS_HOME/librarysystem/stacks/dfilexxx |

3. Move the archive repository files from the original device to the target device.

Note You must delete the archive repository files from the original device. Failure to do so may cause unexpected results.

4. Modify the Device property for the Archive Repository to point to the new target device.

5. Enable logons by setting the Logons Enabled? property in the System object to Yes.

Archive Categories

Like storage categories, archive categories determine where a library system moves the archived document versions. For each library system, you can assign the archive categories unique names appropriate to the offline storage capabilities of that library system. For example, for a library system that archives files onto tape, sends them to a printer, or simply deletes them, you might set up archive categories named: Tape Archive, Print Archive, and Delete Without Backup. For another library system you could assign names such as Six-Month Hold, Old Docs, and Trash.

Whether or not a user can reclaim (to online status) the document versions handled by these archive categories depends on the needs and limitations of the particular library system. You can designate each archive repository as reclaimable or nonreclaimable. A library system also provides an option for the nonreclaimable archive repository to immediately delete versions from specified nonreclaimable archive repositories.

During the installation of a library system, Archive Repository 1 (**dfile001**) is automatically created, for each particular library system, to handle reclaimable files: files that can be brought back online after being archived. An archive category (the default is Trash) is also created and assigned to this reclaimable archive repository.

Note If you want to use Trash as an archive category for nonreclaimable files, you must first create a nonreclaimable archive repository, and then reassign Trash to it.

Storage space requirements for archive categories may be similar to those of storage categories, but it is more likely that many of the versions moving offline will be stored only for a short time, compressed for archiving, or even destroyed immediately.

It will be important for you to determine the archive/reclaim requirements of your users before you add archive categories to the system. Keeping only a few versions of any given document online may seem to be a good tactic, but if you plan to archive most of the older versions, you will need adequate network resources outside of the library system to do so. Likewise, you will need an efficient maintenance schedule to deal with any associated long-term storage requirements.

Many archive categories may be linked to one archive repository, but you cannot link an archive category to more than one archive repository at the same time. Each archive category is flagged as reclaimable or nonreclaimable according to the archive repository that it is currently linked to. Archive categories are linked to archive repositories in the same way that storage categories are linked to storage repositories.

When a document is added to a library system, a default archive category is assigned to it. An authorized user, however, can change this initial value in the Item object.

You can change the assigned archive repository for an archive category at any time. Thereafter, any versions of documents that are moved offline using this archive category will be stored on the newly assigned archive repository. Such a change is often useful in controlling the distribution of files across network disk space.

Care should be taken when assigning a new archive repository, however, because the Reclaimable? property value from the new archive repository will be applied to the archive category. Although the Reclaimable? property is informational only, you will want to avoid the confusion that may be caused by moving reclaimable files onto nonreclaimable archive repositories, and vice versa.

Also, if you add an archive category to a nonreclaimable archive repository, be sure to note whether the archive repository is set up to immediately delete, rather than temporarily store, archived versions. That is, check the value of the Delete Upon Archive? property in the Archive Repository object.

Caution We recommend you name carefully those archive categories that point to archive repositories from which files will be immediately deleted. Names such as Quick Delete or No Reclaim will remind users and yourself that files, once archived in this way, are not reclaimable.

Each archive category has its own object instance with Archive Category (name), Comment, and Assigned Archive Repository properties. You provide values for these properties when you add the archive category to the library system. Most sites begin by adding an initial set of archive categories and add more as needed. Although you can add new archive categories at any time and have as many as you want, you cannot remove or rename them. This restriction is enforced because each existing archive category may have been used in many places throughout the library system.

Manual archiving is a slightly different case, since the user may override the document's archive category for particular versions and assign archive categories that are more appropriate. However, the library system still handles the actual transfer of the versions from storage repositories to archive repositories. The user does not have to be concerned about the disk space used or the actual network locations of the stored versions.

Adding an Archive Category

To add an archive category:

Note For tool-specific instructions on how to add an archive category, refer to the *CS Explorer Online Help*.

1. Navigate to the archive categories for this library system.
2. Add an archive category.
3. Enter property values for the new archive category. Required properties include:

- Archive Category (name)

Double-check the spelling of the archive category name. Once the archive category has been added to the system, you cannot modify the name.

The specified name should clearly indicate to users whether version files assigned to this archive category will be reclaimable once they are moved offline. The archive repository you assign to this archive category will determine whether the files are reclaimable or not and whether they will be immediately deleted as they are archived.

- Assigned Archive Repository

If you do not specify an archive repository, the system will assign Archive Repository 1 by default.

4. Save your changes.

Modifying Archive Categories

To modify an archive category:

Note For tool-specific instructions on how to modify an archive category, refer to the *CS Explorer Online Help*.

1. Navigate to the archive categories for this library system.
2. Select the archive category you want to change.
3. Modify the property values you want to change. Modifiable properties include:

- Assigned Archive Repository
 - Comment
4. Save your changes.

Content Search Repositories

Content Search Managers index files that users want to be able to find later based on the contents of the current versions of those files. Content search functionality is an optional component of a library system.

The Content Search Manager contains the necessary executables and program files to enable use of content searches, as well as an initial content search repository for that service. A content search repository keeps track of the contents of the files that users tag for content search purposes. You can have as many as 64 content search repositories per library system.

To allow files to be indexed for content searches, the library system provides several file filters that will convert the file formats to ASCII. If no filter is available for a file's format, that file will not be indexed. For details on the file formats supported, refer to [“Content Search Filters” on page 315](#).

The library system is configured so that you can avoid having content search repositories filled up with extremely common words that users are not likely to search on (for example, articles such as *a*, *an*, or *the*). Such words are referred to as stopwords, which are text strings in the versions that:

- Are not indexed.
- Cannot be searched for.
- Are counted in the specified distances between terms in phrase searches.

To configure Content Search Stopwords, refer to [“Configuring Content Search Stopwords” on page 320](#).

Caution You can add or delete from this stopword list using any text editor, but you must make these changes before you enable content searches at your site. Otherwise, you will have to reindex all existing indexed versions so that the new stopword list is applied. You must type each stopword on its own line and save the file in text-only format.

When users index files for content searches (or deindex such files), they are actually submitting requests, which are added to a queue. At specified intervals, which you can configure, the Storage Manager starts an indexing daemon. The indexing daemon processes the files in the queue through the Content Search Manager. The daemon will run until it has processed the number of bytes specified in the Content Search Repository object or until it reaches the last file in the queue.

Where you install the Content Search Manager (and its repositories) can affect the amount of traffic on your network, especially in complex library systems having numerous library systems and servers. Versions are indexed in the content search repositories based on their assigned storage categories. Storage categories are, in turn, assigned to particular storage repositories. Therefore, to reduce the flow of network traffic, try to place the Content Search Manager and its repositories on the same server as those storage repositories having storage categories that are most likely to be indexed for content-based searches.

To add a content search repository, you must know the name of the server on which the Content Search Manager is installed. You must also know on which device or drive you want the repository placed.

After you have added the content search repository, you should make sure the following values are set appropriately for your site.

| Setting | Description |
|---------------------|---|
| CS Index | The maximum number of seconds you want the system to wait before attempting to restart the indexing daemon (the default is 300). Refer to “Storage Manager Tab” on page 167 for information about setting this value. |
| Max Repository Size | The maximum allowable size (in kilobytes) of the repository. The installation program initially sets this value to 0 (unlimited). |

By adjusting the indexing interval value (CS Index), you can control how often the indexing daemon runs to process the queue. The best settings for your site will depend on the number and size of files being indexed/deindexed. For example, if you or your users notice that queued indexing requests are not being processed quickly enough, you can reduce the Indexing Interval value so that indexing is done more often.

By specifying a particular value for the Max Repository Size property, you can control the amount of disk space the repository can take up. FileNet recommends that you leave free at least 50% more disk space than the maximum size of each content search repository. This additional disk space is used for temporary files that are created whenever the indexing daemon processes queued indexing actions and updates the size and contents of the repository.

After you add the content search repository, you must also decide which storage categories you want to assign to it. Using CS Explorer, you can assign a content search repository in each affected Storage Category object, as follows:

| Storage Category Object Setting | Description |
|------------------------------------|--|
| Assigned Content Search Repository | The content search repository that will handle the indexed files from this storage category. |

Note The system automatically assigns Content Search Repository 1 to the initial storage category (the default name is Documents). You must assign a content search repository to other existing storage categories. If you add more content search repositories, you can change the assigned content search repository for any storage category.

Once the Content Search Manager is installed and its repositories are created, you can enable and disable content searches system-wide by toggling the value of the Content Search Enabled? property in the System object.

You can also establish settings that will provide a default behavior for indexing, or not indexing, new documents and versions for content searches when they are added to the system. This behavior is controlled by the Default Index Versions? property in User, Group, and System objects.

When you install a Content Search Manager, the library system installation program creates a content directory on the installation device you specify (for example, under UNIX, `$IDMDS_HOME/librarysystem/content`). Required content search executables and program files are added to the bin and data subdirectories in this path. A third subdirectory, **index**, contains the content search repository.

If you place the content search repository on the same device as the Content Search Manager, the index subdirectory will appear with the bin or data subdirectories. Otherwise, it will appear in its own path under a separate content directory. It will contain the content search repository files and the two log files described in the following list (where x is the content search repository ID):

| Content Search Logfile | Description |
|------------------------|---|
| index0xx.log | Tracks the size of the content search repository and the volume of processed data. |
| csi0xx.log | Provides useful information about content search indexing operations on the server (such as, when indexing has started and stopped, and indexing failures). |

Note The library system allows these indexing logs to grow to as large as 250 kilobytes. When this limit is reached, the log files are emptied and the contents are saved to **index0xx.lsv** and **csi0xx.lsv**. These .lsv files are overwritten each time this process occurs. If you are having problems with content search software, the information in the **csi0xx.log** (or **csi0xx.lsv**) file may provide valuable information to you and FileNet Technical Support.

Adding a Content Search Repository

To add a content search repository:

Note For tool-specific instructions on how to add a content search repository, refer to the *CS Explorer Online Help*.

1. Make certain that a Content Search Manager has been installed on the same server on which you plan to add the content search repository.

For best performance, we recommend adding only one content search repository to each Content Search Manager.

No more than 64 content search repositories can be added to a single system.

2. Navigate to the content search repositories for this library system.
3. Add a content search repository.
4. Enter property values for the new content search repository. Required properties include:
 - Server
 - Device

You cannot add a content search repository to removable media such as a floppy disk.

5. Save your changes.
6. If a message indicates that disk space is not available for the new content search repository, you must do one or more of the following:
 - Delete (or store offline) some of the files or directories from the current disk.
 - Obtain a larger disk.
 - Add the Content Search Manager and repository to a different library system server (new or existing).
7. In one or more Storage Category objects, assign the content search repositories you have added. Versions in these storage categories, if indexed for content searches, will be tracked in the assigned content search repository.

8. Stop and restart the Content Search Manager on the library system server to which you added the content search repository. Refer to “Controlling Library System Services” on page 198 for information about stopping, starting, and checking the status of services.

To determine whether the content search repository is correctly attached to its Verity K2 Server, and to attach it if it isn't, run the utility CSK2CFG. Refer to “CSK2CFG” on page 386 for more information.

Modifying a Content Search Repository

To modify a content search repository:

Note For tool-specific instructions on how to modify a content search repository, refer to the *CS Explorer Online Help*.

1. Navigate to the content search repositories for this library system.
2. Select the content search repository you want to change.
3. Modify the property values you want to change. Modifiable properties include:
 - Bytes of Working File
 - Bytes of Working Memory
 - Device (Refer to “Moving a Content Search Repository” on page 256 for instructions on how to modify this property.)
 - Hour of Day to Merge
 - Logging Search Server SQL?
 - Max Immediate File Size
 - Max Repository Size
 - Network Path
 - Server
4. Save your changes.

Moving a Content Search Repository

To move a content search repository:

Note For tool-specific instructions on how to move a content search repository, refer to the *CS Explorer Online Help*.

1. Put the library system in single-user mode. A library system is in single-user mode when:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.

- On the target device, recreate the directory path that contains the content search repository files. The content search files can be found in:

| OS Platform | Location |
|-------------|---|
| Windows | <i>Drive:\filenet\LibrarySystem\Content\Index</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/content/index</i> |

- Move the content search repository files from the original device to the target device. The original device may contain files for several content search repositories, so make sure you only select the files for the content search repository you want to move (for example, **index002.***).

Note You must delete the content search repository files from the original device. Failure to do so may cause unexpected results.

- Modify the Device property for the content search repository to point to the new target device.
- Run the CSK2CFG utility to register and attach the Verity components of the library system you are restoring.
- Run the REPCHECK utility to verify that the logical links between the content search repository and the property repository. Respond to any discrepancies found in the REPCHECK utility report log. Refer to “REPCHECK” on page 439 for instructions.
- Enable logons by setting the Logons Enabled? property in the System object to Yes.
- Stop and restart the Content Search Manager service. Refer to “Controlling Library System Services” on page 198 for information about stopping, starting, and checking the status of services.

To determine whether the content search repository is correctly attached to its Verity K2 Server, and to attach it if it isn't, run the utility CSK2CFG. Refer to “CSK2CFG” on page 386 for more information.

Custom Properties

All documents that are added to a library system are described with an extensive set of properties that allow users to quickly identify, locate, and retrieve them. Additional properties can be added to Item, Version, and Checkout objects and assigned customized names and purposes. These additional properties are referred to as custom properties.

Custom properties can be single-value or multivalued. Single-value custom properties can accept a *single* string, integer, or date and time value. Multivalued custom properties can accept *multiple* values of a *single* data type (either string, integer, or date and time). Up to 100 single-value and 10 multivalued custom properties can be created per Item, Version, or Checkout object.

Internal Custom Properties

The following internal custom properties are created during the library system installation or upgrade. They are for internal use only. DO NOT open them, edit them, or attempt to delete them. Any changes to these internal custom properties can result in system failure.

- Document Class
- IdmAnnotation
- IdmCDBehaviorID
- IdmFilterQuery
- IdmPublish
- IdmSearch
- IdmVWVersion (used only by eProcess)

Characteristics of a Custom Property

You can set up each custom property to have the following characteristics:

- (Required) Have a name customized for the particular system.
- (Required) Have a parent object type (Item, Version, or Checkout).
- (Required) Accept a string, integer, or date and time value.
- (Required) Have a set maximum value length.
- (Required) Be single-value or multivalued.
- (Optional) Have a default value.
- (Optional) Be a required user entry.
- (Optional) Have different display names.
- (Optional) Have a quick list enabled to allow faster display of selectable values.

- (Optional) Have indexing of the custom property's values enabled (which may improve search performance if the quick list is enabled or the custom property has few distinct values).
- (Optional) Assign access rights to the custom property.

The other properties in the Custom Property object (for example, Date Added and Comment) are provided by the library system. Before adding your own custom properties, you should review the names and purposes of these properties to avoid creating duplicates.

Adding Custom Properties

Note For tool-specific instructions on how to add a Custom Property, refer to the *CS Explorer Online Help*.

Caution Several internal custom properties are created during Content Services installation or upgrade. Although these custom properties are displayed in CS Explorer, you must not attempt to delete them or modify them in any way. Refer to “[Internal Custom Properties](#)” on page 258 for details.

Consider the following suggestions and recommendations when adding custom properties:

- **Check your user interface requirements first.** Be aware that some interfaces to a library system will create a number of custom properties as they are installed. Before you add custom properties, check the interface requirements and the property names assigned.
- **Pre-determine departmental or system-wide needs.** Keep in mind that the names of the custom properties are system-wide and apply to the documents added by all users in all groups. A unique name that makes sense to one department may not be applicable in another department. A custom property can be defined as an optional entry, but if the company has many such dissimilar departments and property requirements, it may be a good idea to have more than one library system. Each library system can then have its own set of custom properties. However, if you are installing FileNet Replication, you will likely need some identical custom properties in each participating library system, so that the replicated documents can be commonly identified during a search.
- **Determine display order ahead of time.** Custom properties will be displayed in Item, Version, and Checkout object properties in CS Explorer in the order in which you add them. If you want the custom properties to be displayed in a particular order, be sure to plan ahead before you begin adding them to any library systems at your site. This could mean that you add the custom properties from the most important (or most frequently used) property to the least important property. But it could also mean that you add the properties in a logical order (for example, Name, Street Address, City, State, Zip Code).
- **Determine maximum value lengths.** For all multivalue custom properties whose property type is string, you can set a maximum value length of 1 to 200 characters. For single-value and multivalue custom properties whose property type is integer, integer string, or date/time, the value length is fixed (see the following table).

A maximum of 100 single-value custom properties can be defined.

(For Oracle only) There is no byte limit on sum of the lengths of all single-value custom properties.

(For Microsoft SQL Server only) For each parent object type, there is a byte limit on the sum of the lengths of all single-value custom properties of that parent type:

- Item parent = 7200 bytes maximum
- Version parent = 5400 bytes maximum

- Version Checkout parent = 6800 bytes maximum

You may also want to keep some space available for future use. The following table shows the lengths of property values for both single-value and multivalued custom properties. Only string properties can have their value lengths increased. (Value lengths cannot be decreased.) Also, the maximum length for a string property value of a single-value custom property differs from that of a multivalued custom property.

Note For each multivalued custom property, you can enter as many values as your database engine and system resources will bear. Content Services does not impose an inherent limit on the combined length of multiple values.

| Property Type | Length | Size |
|-----------------------|---------------------|---|
| string (single-value) | 1 to 254 characters | 254 bytes (defined by Property Length property of Custom Property object) |
| string (multivalued) | 1 to 200 characters | 200 bytes (defined by Property Length property of Custom Property object) |
| integer | 10 digits | 4 bytes |
| integer string | 9 digits | 10 bytes |
| date/time | 32 characters | 10 bytes |

Note Integer values can contain a leading '+' or '-' sign followed by at most 10 digits. The '+' is optional for positive integers. Thus, an unsigned integer value contains at most 10 characters; a signed integer contains at most 11 characters (one character to accommodate the sign). Leading zeros are not allowed. The valid range for integer custom properties is -2,147,483,648 to +2,147,483,647. Entering a value outside this range will result in an error.

Integer string values can contain an optional leading '+' or '-' followed by at most nine digits. Leading zeros are allowed. For example, if you set the value of an integer custom property to '123' and the value of an integer string custom property to '00123', the database will store the values as 123 for the integer and 00123 for the integer string.

- **Assign meaningful characteristics.** Before deciding on names for the custom properties, analyze how the future users of the library system currently look for the objects with which they work. For example, if they search for file folders by an invoice number under a particular client name, then it would make sense to provide properties called Invoice Number and Client Name. Although you will be able to rename custom properties, you will not be able to delete them or change several of their characteristics (for example, to change parent object or property type).
- **Keep properties for future use.** You can add no more than 100 single-value and 10 multivalued custom properties per parent object. Because it is not possible to delete a custom property from the system, it is a good idea to keep several custom properties in reserve to allow for evolving user needs.

Note Most sites add a starter set right away and add more from time to time, as dictated by the evolving needs of the users. Documents that are added before a particular custom property has been added will have no value for that property in their Item or Version objects. But, a user with Owner or Admin access rights to the Item or Version object can supply a value for the new property at any time.

- **Assign access rights carefully.** Be careful about assigning access rights to the custom properties, especially if you assign access rights of None. We recommend against assigning access rights of None to a required-entry custom property, especially for users who can add documents. For example, if you were to assign None access rights to the Owner Privileges property in the Custom Property object, any user adding a document would be required to enter a value for the custom property but could not see that property thereafter on the Item object.
- **Index a property.** We recommend you allow the values for a custom property to be indexed if either of the following is true:
 - You designate a quick list of values for the property.
 - You commonly perform searches that use the operators =, <, >, or “starts with.”

Otherwise, you should change the Indexed? property to No. Be aware, too, that system performance can be adversely affected if more than three or four custom properties are indexed. You will have to experiment at your site to determine the optimum number of indexes for the operational demands placed on your systems.

- **Assign display names for special use.** We strongly recommend that you assign custom property names that are unique from any other custom or system-generated property names. You may have reason at your site to assign one or more display names based on support for different national languages provided by particular user interfaces. However, you may also want to add display names that mask the actual custom property names that you use as administrative identifiers. For example, you might add a custom property named Finance Subject, (so that you know its exact nature) but want the name displayed simply as Subject to the users of the library system that controls the Finance Department's data.

Modifying Custom Properties

For tool-specific instructions on how to modify a custom property, refer to the *CS Explorer Online Help*.

Caution Several internal custom properties are created during library system installation or upgrade. Although these custom properties are displayed in CS Explorer, you must not attempt to delete them or modify them in any way. In particular, the Document Class property is required for document class-based functionality in eProcess Administrator, such as automatic launch; renaming or removing the property will disrupt this functionality. Refer to “[Internal Custom Properties](#)” on page 258 for further details.

To modify custom properties:

1. Navigate to the custom properties for this library system.
2. Select the custom property you want to change.
3. Modify the property values you want to change. Modifiable properties include:
 - Author Privileges
 - Comment
 - Indexed?
 - Owner Privileges
 - Property Length

Note You cannot decrease the length of a custom property; you can only increase it.

- Property Name
- Viewer Privileges
- Quick List?

You can also modify display names from the Custom Property object.

4. Save your changes.

User Objects

Before someone can open a session and begin using a library system, his or her name must be added to the internal list of users maintained by the system. Generally, an administrator adds an initial number of users and groups, and then, may delegate the authority to add new users and groups to one or more of these users.

Notes

- Administrators always have the right to add users, groups, and documents to the library system, regardless of the property values in their User object.
- When using CS Bulk Process to add users to a library system, do not add more than 500 users within one bulk process.

Every user has a User object that contains descriptive properties. Some of these properties provide default values. For example, the Default Group property can be set to the group in which the user does the most work. Then, each time the user opens a session, his or her active group will be set to this group automatically.

Other properties allow or deny the user certain privileges, such as adding documents to the library system.

Working with User Objects

Note For tool-specific instructions on adding, modifying, renaming, or deleting User objects, refer to the *CS Configuration Tool User Import Tab*.

Do not use the Euro character or other non-alphanumeric or special characters, such as '~ ! @ # \$ % ^ & * () _ + - = [] \ { } | ; ` : " , . / < > ? in user names. (FN DTS00067196)

Also see “[Importing Users and Groups](#)” on page 272 for information about importing users from your network domain or LDAP directory to your library system.

Adding a User Object

When you add a user to a library system, the system enters your name as the value of the Added By User property in the new User object. This user (you) is also added to the new User object access list with Admin access rights. (The name of the new user is added to the access list with Owner access rights.)

Giving Admin access rights to the person who adds the user allows that person to maintain all user-modifiable properties in the new User object. The following properties are modifiable only by users with these Admin access rights:

- Add Items?
- Add Users?
- Add Groups?
- Enabled?
- Password
- Primary Storage Manager

- User Name
- Secure Searches?

After you add a new user to the library system, you will want to make the new user a member of one or more groups. This must be done from the Group object in the particular library system. You can then use FileNet Replication Services to replicate user and group objects to other participating library systems.

Renaming a User Object

An administrator can change the User Name of any library system user except Admin. The library system will not allow you to:

- Rename the Admin user
- Assign an existing user name
- Assign blank name value to a user

Deleting a User Object

When a user in the library system leaves your company or transfers to a different site, you may want to delete the associated User object from the library system. With CS Explorer, an administrator can delete a User object from a library system.

When you delete a User object, that object is removed from the library system. Also, the user is deleted from all access lists and group membership lists. The user's name, however, will remain in single-value properties, such as the Added By User property in Item objects, which ensures that this value can still be used in searches.

Assigning and Modifying Values in a User Object

The following topics describe some of the most important properties in the User object. For additional property information, refer to *CS Explorer Online Help*. These properties can be grouped together into two categories: properties unique to a User object and properties common to User, Group, and System objects.

Properties Unique to the User Object

The properties described below are unique to the User object and cannot be inherited from a Group object or the System object.

Enabling a User to Add Items, Users, or Groups

The Add Items, Add Users, and Add Groups properties determine what a particular user can add to the library system. (The exception to this rule is that all administrators can add users, groups, and documents to the library system even when the associated properties in their User objects are set to No.)

By default, the library system withholds from all new users the authority to add documents, users, and groups. However, you will probably want to change the Add Items value to Yes when you add a new user to a particular library system. Moreover, a user cannot add annotations to a document in the library system if the Add Items value of his User object is set to No. Withholding the authority to add and annotate documents typically makes sense only to support an online service where a small number of users create and maintain information for a larger group of users who only need to view the information in the library system.

Note A user can apply Annotations to documents in the library system only if the “Add Items?” property is set to Yes in his/her User object. (However, the ability to modify the content of the item is based on the Access Control List, not the Add Items? property value.)

Since the library system allows an unlimited number of administrators, the need to delegate the authority to add users and groups to people who are not administrators may not be obvious. In fact, at some sites, the authority to add users and groups is not delegated beyond simply adding more administrators. At other sites, however, there may be only a part-time administrator, and it may be inappropriate to have more than one individual with this level of access rights. In this case, it makes sense to delegate the authority to add users and groups to several users who are not administrators.

Enabling/Disabling a User's Access to the System

The Enabled? property determines whether a user can access the library system (whether he or she is currently allowed to open a session). When the user is first added to a particular library system, the value for this property is set to Yes by default. However, if you want to completely restrict a user's access to the library system (typically, for an extended period of time, such as during a long leave of absence), you can simply set this Enabled? value to No.

If you are storing particularly sensitive information, you can use this feature to protect not only that information, but also the user's own data and access integrity. Let's say a site has automated the library system logons (for example, basing them on network workstation ID's) so that normal library system logon procedures are bypassed. At such a site, setting this Enabled? property to No would prevent unauthorized personnel from gaining access to the library system at a user's vacant workstation.

Setting Logon Values

The Network Logon Name property in the User object plays a key role in automating the logon process. For details, refer to the *CS Explorer Online Help*.

Modifying the User Object Access List

As mentioned earlier, each time a new user is added to the library system, the person who added the user is granted Admin access rights to the object and the new user is granted Owner access rights. This person can then add, change, or remove any entries from the access list.

Properties Common to User, Group, and System Objects

The properties described below are common to the User, Group, and System Objects and can be inherited. Refer to “[Understanding Default Values and the Inheritance Path](#)” on page 188.

Setting User Default Values

While working with a User object, you need to be aware of how the library system handles default values. The User, Group, and System objects each have a Default Archive Category, Default Storage Category, Default Checkout File Extension, Default Checkout File Name, Default Item Access List, and Default Index Versions property. However, even though these three objects have properties with the same names, the properties are still completely independent of each other and can be given different values in each object.

Note If you delete or disable a user's default group and that user opens a session, the active group will be the General Users group rather than the group specified in the Default Group property.

When a user adds a document to a library system, the system uses these default properties to supply some of the initial values in the new Item object.

Controlling the User's Search Results

The following two properties control what results are returned for users' searches.

| Property | Description |
|------------------|--|
| Access Domain | The Access Domain property controls whether a user's access rights to search results are based on the access rights of his or her active group or on the highest access rights of any of the groups to which the user belongs. The Access Domain property appears in User, Group, and System objects. |
| Secure Searches? | The Secure Searches? property controls whether the user must have at least Viewer access rights to a document (or object) to have that document (or object) included in the results of a search that he or she runs. Only users with Admin access to the User object can modify this property. The Secure Searches? property appears in User, Group, and System objects. |

Specifying the User's Primary Storage Manager

When a user access an item, the Primary Storage Manager property specifies the Storage Manager that accesses the database information about the item. By providing a value for this property at the User object level, this Storage Manager processes the user's requests. Assigning different Primary Storage Managers to different users is one way to load-balance activity across servers.

Assigning a Different Display Language

The Display Language property controls the user's display language for property names, field names, and messages (posted to the User object or the Administrators Group object). Like the default properties mentioned in a previous topic, the Display Language property appears in User, Group, and System objects.

Group Objects

A library system predefines two special groups: General Users and Administrators. When a user is added to a library system, he or she is automatically made a member of the General Users group. Therefore, all users have access to any document that grants the General Users group some kind of access rights. Of course, the General Users group can be granted access rights of None or even deleted from the document's access list altogether. The General Users group's access rights in most library systems are restricted to Viewer.

The users who are members of the Administrators group are referred to as administrators. When they designate this group as their active group, these users have universal Admin access rights to all objects in the affected library system. In addition, administrators always have the right to add users, groups, and documents to the library system (even if the corresponding properties in their User objects are set to No), and administrators can always cancel another user's checkout.

Before adding any groups to the library system, observe the people who will be using the system and the way they are presently working together. Determine natural groupings of these people based on their common access requirements to the files (or information) that will eventually be added to the library system.

You may be able to use the formal organization of the company to identify potential groups and control access with them. But with a library system, you also have the opportunity to establish additional groups that do not exist in the actual organization in any form. This capability allows you to create groups for people who work together informally across the formal organizational boundaries.

For example, suppose that a particular library system was company wide. In addition to the various departmental based groups, there could be groups for the Employee Recreational Committee, the Capital Equipment Approval Board, and so on.

The existence of groups makes access control within the library system straightforward. There can be many groups in the library system. However, no group may be a member of another group. A user may belong to more than one group.

When logging on, a user's default group becomes his active group. When working in the library system (depending on the client interface to the library), a user may designate another group (that he is a member of) as his active group. The active group determines what documents can be accessed by the user. A user may designate a different group as his or her active group at any time during a session, as long as that user is already a member of that group.

One advantage to setting up the library system with groups that match the way the users actually work together is that the users can readily relate to the structure. They do not have to learn an entirely new way of working.

Before adding any groups to the library system, observe the people who will be using the system and the way they are presently working together. Determine natural groupings of these people based on their common access requirements to the files (or information) that will eventually be added to the library system.

Working with Group Objects

Note For tool-specific instructions on adding, modifying, renaming, or deleting Group objects, refer to the *CS Explorer Online Help*.

Do not use the Euro character or other non-alphanumeric or special characters, such as ' ~ ! @ # \$ % ^ & * () _ + - = [] \ { } | ; ` : " , . / < > ? in group names. (FN DTS00067196)

Also see “[Importing Users and Groups](#)” on page 272 for information about importing groups from your network domain or LDAP directory to your library system.

Adding a Group Object

When you add a new group to the library system, your name becomes the value for the Added By User property in the new Group object, and you are granted Owner access rights in the access list of the new object. Of course, as an administrator, you also have Admin access rights to this object.

The following properties can be modified only by a user with Admin access rights to the Group object:

- Enabled?
- Group Name
- Primary Storage Manager
- Secure Searches?

You can use FileNet Replication Services to replicate group objects to other participating library systems.

Renaming a Group

An administrator can rename the Group object assigned to any group that has been added to the library system. Content Services will not allow you to:

- Rename the General Users group
- Rename the Administrators group
- Assign an existing group name
- Assign a blank name value to a group

Deleting a Group Object

Undoubtedly, there will be times when groups no longer have viable functions in your organization. You will probably want to remove these groups from the library system at some point, but without removing all references to them on documents that group members added to the library system or on versions of documents that these members may have created during the course of their work.

An administrator can delete a Group object from the library system while retaining all existing references to that group. When you delete a group, the associated Group object is removed from the library system. Also, the group is deleted from all access lists. The group's name, however, will remain in single-value properties, such as the Added By User's Group property in Item objects, which ensures that this value can still be used in searches.

Assigning and Modifying Values in a Group Object

At some sites, Group objects are administered by managers or lead members of a work group or company organization. The library system administrator may have little reason to be involved in the day-to-day modifying and updating of the Group objects for a particular library system. Depending on how your library system is set up, the person who adds the group to the system may or may not be the person who updates the Group object.

The following topics describe some of the most important properties in the Group object. For additional property information, refer to *CS Explorer Online Help*. These properties can be grouped together into two categories: properties unique to a Group object and properties common to User, Group, and System objects.

Properties Unique to the Group Object

The properties described below are unique to the Group object and cannot be inherited from the System object.

Enabling/Disabling a Group's Status

The Enabled? property determines a group's status (that is, whether users are currently allowed to open a session under or become an active member of this group). When a group is first added to the library system, the value for this property is set to Yes by default. However, if you want to completely restrict a group's access to the library system for an extended period of time, you can simply set this Enabled? value to No.

Note If a group is disabled and a user has this group specified as his or her default group, the library system will still allow the user to open a session, but only as a member of the General Users group. Some user interfaces to the library system may or may not work this way.

If you are storing particularly sensitive information, you can use this feature to protect not only that information, but also the data and access integrity of the group's members. For example, you could choose to disable a group for a period of time to investigate breaches of security or to close down the activities of a group that has been temporarily moved off site.

Modifying the Group Object Access List

Each time a new group is added to the library system, the user who adds the group is granted Owner access rights to the object and the General Users group is granted Viewer access rights. When adding the group, you can add, change, or remove entries from the access list.

Properties Common to User, Group, and System Objects

The properties described below are common to the User, Group, and System Objects and can be inherited. Refer to [“Understanding Default Values and the Inheritance Path”](#) on page 188.

Setting Group Default Values

While working in the Group object, you need to be aware of how the library system handles default values. The User, Group, and System objects each have a Default Archive Category, Default Storage Category, Default Checkout File Extension, Default Checkout File Name, Default Item Access List, and Default Index Versions? property. However, even though these three objects have properties with the same names, the properties are still completely independent of each other and can be given different values in each object.

When a user adds a document to a library system, the system uses these default properties to supply some of the initial values in the new Item object. As a library system administrator, you should probably not alter any of the values the group leader has assigned to these default properties in the associated Group object.

Controlling the Group's Search Results

The following two properties control what results are returned for group members' searches:

| Property | Description |
|------------------|---|
| Access Domain | The Access Domain property controls whether a group member's access rights to search results are based on the access rights of his or her active group or on the highest access rights of any of the groups to which the user belongs. The Access Domain property appears in User, Group, and System objects. |
| Secure Searches? | The Secure Searches? property controls whether group members must have at least Viewer access rights to a document (or object) to have that document (or object) included in the results of a search that they run. Only users with Admin access rights to the Group object can modify this property. The Secure Searches? property appears in User, Group, and System objects. |

Specifying a Group's Primary Storage Manager

When a user access an item, the Primary Storage Manager property specifies the Storage Manager that accesses the database information about the item. By providing a value for this property at the Group object level, this Storage Manager processes the user's requests. Assigning different Primary Storage Managers to different groups is one way to load-balance activity across servers.

Assigning a Different Display Language

The Display Language property controls the language used for displaying property names, field names, and messages (posted to the User object or the Administrators Group object). The Display Language property appears in User, Group, and System objects.

Setting Up or Modifying the Administrator's Group

Any user can be an administrator. The only requirement is membership in the Administrators group. If you are a member of the Administrators group and the Administrators group is your active group, then you have Admin access rights to any object, including the System object.

Typically, the initial set of administrators are in a position within a company to understand what the users need. They understand the flow of work between users, the work-group structures, and the flow of work between work groups. They know what kind of information needs to be shared. And, as a general rule, they are proficient users of other software tools.

There are two reasons to add members to the Administrators group:

- The workload demands it.
- You need backup during absences.

Caution The number of library system administrators is best kept to the minimum required. The data within the library system is a valuable company resource. Its integrity and security are difficult to maintain if there are too many administrators or if care is not used in their selection. Also, most sites remove the user named Admin from the Administrators group after they have added other users to the Administrators group.

Special Administrative Groups

You can create two groups, zUserGroupAdmin and zPasswordAdmin, whose member users have limited administrative rights. The users in zUserGroupAdmin can add, delete, and modify CS users and groups; the users in the zPasswordAdmin group can modify passwords of CS users and groups. However, no user in either group will be able to do any of the following:

- Delete or modify the initial Admin user or other “special” users in CS—specifically, the DirSync, ReplDaemon, and WorkflowDaemon users.
- Delete, modify, or add members to the initial Administrators group, the zUserGroupAdmin, or the zPasswordAdmin group.

Assign members to the zUserGroupAdmin group with caution. Remember, they can make changes to *all* modifiable properties in all User and Group objects, which means they can, among other things, do the following (except to the initial Admin user and Administrator group):

- Delete CS users and groups, even administrators.
- Enable/disable CS users, even administrators.
- Change the Access Domain property for any user, even administrators.
- Change the Access List or Default Item Access List entries for any user or group, even administrators or their groups.
- Enable/disable users' rights to add documents, users, and groups.
- Enable/disable secure searches (which controls access to documents and their properties).
- Change defaults for storage category, archive category, and checkout/copy path, which could affect specified storage policies and loads.
- Assign members to the zUserGroupAdmin group with caution. Remember, they can change even administrators' passwords.

Note You can always use the super-user Admin account in CS to undo changes made by members of these quasi-administrative groups, assuming you or other administrators in your organization have ready access to the logon credentials for the Admin user.

Importing Users and Groups

Content Services provides User Import tools that perform one-way synchronization of your library-system users and groups with the Windows network-domain users and groups or with organizational units (users and groups) in an LDAP directory.

These User Import tools are controlled by settings in two tabs of the Content Services Configuration application. To import users/groups from a Windows domain, use the WinNT/Win2000 User Import tab; to import them from an LDAP directory, use the LDAP User Import tab. These tabs are available only on the initial storage server, not on additional storage servers.

The User Import tools that accomplish the synchronization are implemented as daemons that are enabled by the settings in the User Import tabs. When enabled and launched, a daemon performs ongoing one-way synchronization of group and user information from selected Windows network domains or an LDAP directory to a selected library-system property repository. The daemons do not propagate changes from the property repository back to the network domains or LDAP directory.

Note At most 65,536 groups can be imported with the User Import tools.

A User Import daemon identifies users and groups (from a Windows NT domain or LDAP directory) to be added to or synchronized with the users and groups in a library system. Whether the addition/synchronization takes place or not depends on properties of the Group and User Objects, as described in “[User/Group Synchronization Criteria](#)” on page 292.

For information on configuring and performing user/group import from Windows network domains, refer to “[Using Windows User Import](#)” on page 273.

For information on configuring and performing user/group import by querying an LDAP directory or specifying an XML file of users/groups, refer to “[Using LDAP User Import](#)” on page 277 and “[Using XML User Import](#)” on page 281.

Note Users imported from a Windows NT domain into a Content Services library system by the Windows User Import daemon may log onto the library without a password, from IDM Desktop or Web Services, after logging onto the NT domain. This “unified” form of logon also works for users imported from a Novell NDS directory, but only when logging on from IDM Desktop.

Using Windows User Import

The following is the recommended procedure for using the Windows User Import tool. This procedure requires that you generate and review a preview transaction file before you import users and groups to your library system. The procedure also recommends that you disable the tool after you have finished importing users and groups.

While it is possible to have the Windows User Import tool continuously synchronize your library system and Windows domain users and groups, before you pursue such a strategy, you should thoroughly review this documentation to understand the implications and risks.

The Windows User Import synchronization is accomplished by a daemon, **dirsync.exe**, running on the initial storage server under a selected library system's Storage Manager. And the Storage Manager services need to be run as a domain account. Otherwise, User Import will not import the users in the selected groups.

If you use the WinNT/Win2000 User Import tab to govern the importing of users into a library system from Active Directory, your Windows domain must be mixed mode. Also, you must set (via IDM Configure) the “Use operating system credentials for logon” preference to “Yes” on any machine where you want to use IDM Desktop/Web Services to log on to the library.

Caution Because the Windows User Import daemon can affect the accessibility and security of the data in your library systems, unless you are very familiar with Windows NT domains we recommend you use this tool only in single domain or simple multi-domain environment. If you wish to use it in a more complex domain configuration, we suggest you first run it thoroughly with your configuration in a test environment so that you understand how the User Import features can be used to dynamically add and delete users and groups.

To configure and start or stop the Windows User Import tool:

1. Open the Windows Control Panel and start the CS Configuration Manager.
2. Select the Installation tab.
3. In the Library System list box, select the name of the library system for which you want to set up Windows domain synchronization.
4. Select the WinNT/Win2000 option button as the Directory Server Type. However, if you previously selected another Directory Server Type, stop the Storage Manager, select the Directory Server Type, and then restart the Storage Manager.

Note If you previously selected a different Directory Server Type, a message will pop up to inform you that the Directory Server **group.log** settings will be lost. Click OK to continue or Cancel to retain your earlier Directory Server Type.

5. Select the WinNT/Win2000 User Import tab.

Note If you select the LDAP User Import tab or the XML User Import tab instead of the WinNT/Win2000 User Import tab, a message will pop up to remind you to use the tab corresponding to the Directory Server Type you chose in step 4.

If you want to change your selected Directory Server Type, continue at step 2 of “[Using LDAP User Import](#)” on page 277 or “[Using XML User Import](#)” on page 281, depending on your new selection.

6. In the Interval field, specify how often (in seconds) you want the daemon to run. The default is to run hourly (every 3600 seconds). The ideal value for this field will depend on how you plan to use the tool. If you plan to

synchronize users and groups and then disable the tool, as recommended by this procedure, you may want to specify a shorter interval (for example, 360 seconds).

If you wish, you can enter an Offset value (in seconds) to offset the daemon's starting time from other processes that run under the Storage Manager. The default offset is 25 minutes (1500 seconds). Refer to [“Storage Manager Tab” on page 167](#) for information about the other Storage Manager processes with configurable offsets.

7. You can specify any valid path and file name for the log this tool produces. The default log file is ***Drive:\filenet\LibrarySystem\dirsync.log***.

If the specified file already exists, new logging information will be appended to it. Only changes made since the previous execution are logged.

8. Enter a maximum size (in kilobytes) for the log file.
9. Select the Windows domain containing the groups and users you want to import into your library system. If the domain you want does not appear in the dropdown list, try typing in the domain name. If the selected domain's groups do not appear in the Available Groups box, try left-clicking in the box.

Notes If you create a group on the primary domain controller and then rename it, the original name, not the new name, will appear in the Available Groups list box of the Windows User Import tab.

Only Security groups with Global scope will appear in the Available Groups list box. No other groups may be selected. Refer to [“User Import Usage Guidelines” on page 285](#) for more information.

The list of domains that can be displayed here is truncated at the 62nd domain. However, you can enter any valid Windows domain name manually.

10. In the Available Groups box, select the entities you want to import and click the right arrow to move them to the Selected Groups box. If the Selected box already contains the names of groups you have previously imported, do not remove them from the Selected box unless you want those groups deleted from the library system.

Caution Subject to the settings of the *Imported Group?* and *Imported User?* properties described in [“User/Group Synchronization Criteria” on page 292](#), the Windows User Import tool will delete from the library system any domain groups that were imported using this tool and then removed from the Selected Group list for a subsequent execution of this tool.

11. Check the “Enable user import” box. Note that the daemon will not be started until you click Apply or OK.

Note The Storage Manager must be running for the daemon to execute.

12. In the Selected Groups box:

- a. To control the order in which groups are processed, you can highlight a group name and use the arrow buttons to move the group up or down the list.

Note If you have multiple groups selected, the order of processing may be important. If the Windows User Import tool assigns a default library-system group (see conditions in [“User/Group Synchronization Criteria” on page 292](#)) it will assign the group created for the first domain group processed.

13. Check the “Create preview file only” box to generate a preview transaction file without making any changes to the library system.

You can enter any valid pathname for the preview file the Windows User Import tool produces. By default, the preview file is written to *Drive:\filenet\LibrarySystem\preview.log*. Refer to “[User Import Preview Transaction File](#)” on page 307 for information about the file format and contents.

14. Enter a valid pathname for the group-list file the Windows User Import tool produces. This is the file that contains all the groups that are to be imported into the library system. By default, the group-list file name is **group.log**.
15. Once you have all of the options set the way you want them, click Apply. The preview log file will be generated when the daemon is run (as determined by the value you entered in the Run interval field).
16. After the Windows User Import tool has executed, review the preview transaction file to verify the changes that will be made to the library system database. Refer to “[User Import Preview Transaction File](#)” on page 307 for a description of the file's content and format.
17. If you are satisfied with the changes outlined in the preview transaction log file, clear the “Create preview file only” box and click Apply. The next time the daemon runs, your library system users and groups will be synchronized with the selected Windows NT domain users and groups.
18. After the daemon has run and your library system has been updated, clear the “Enable user import” check box and click Apply.
19. Review the **dirsync.log** file to verify that the desired users and groups were added to your library system.
20. Newly added users may have to immediately change their library system user password. Refer to “[Passwords for Imported Users](#)” on page 284 for more information.

Note WinNT/Win2000 User Import supports only mixed mode. In this mode only global security groups are imported.

Installing an iPlanet LDAP Server

To import users and groups from an LDAP directory into a Content Services library system, you need to install an LDAP server, which must be iPlanet Directory Server 5.0 (or higher) running on Sun Solaris 8, Windows 2000 or Windows 2003, or HP-UX.

For installation and administration instructions, refer to *iPlanet Directory Server Installation Guide* and *iPlanet Directory Server Administrator's Guide* on the iPlanet installation CD.

iPlanet Directory Server contains several related software components. The core of the iPlanet LDAP directory application is the Directory Server. The iPlanet Console (which can be installed as a standalone application on any machine) provides a common user interface to the Directory Server, as well as to other iPlanet server products. The Administration Server serves as the front-end for an iPlanet server and is the interface between Directory Server and iPlanet Console.

The iPlanet installation procedure will guide you through the installation of each of these components. Although there are four different installation methods available to you, it is recommended that you follow the Typical Installation or the Custom Installation method.

The Typical Installation method allows you to create an initially empty Directory Server. The Custom Installation is similar to the Typical Installation but also allows you to import an LDIF file to initialize the directory database that is created by default.

Note During iPlanet installation, be sure to specify passwords of at least eight characters for the *admin* user, as well as other users, even though the installation procedure allows you set passwords with fewer than eight characters. These longer passwords (at least eight characters) are required for proper functioning of the LDAP_EXP and related iPlanet utilities.

Before starting the iPlanet installation, be sure that your machine satisfies the hardware and software requirements for iPlanet Directory Server. Refer to “Chapter 2 Computer System Requirements” of the documentation on the installation CD. Also, be sure to check [iPlanet](#) for the latest release notes and required iPlanet patches.

Using LDAP User Import

The following is the recommended procedure for using the LDAP User Import tool. While it is possible to have the LDAP User Import tool continuously synchronize your library system and LDAP-directory users and groups, before you pursue such a strategy, you should thoroughly review this documentation to understand the implications and risks. You may want to disable LDAP User Import after you have finished importing users and groups.

Before you can import users and groups from an LDAP directory, you must set up an LDAP server. The server must be one of the following:

- Windows Active Directory (native mode only)
- iPlanet Directory Server 5.0 (or higher) running on Sun Solaris 8, Windows 2000 or Windows 2003, or HP-UX (refer to “[Installing an iPlanet LDAP Server](#)” on page 276 for installation information)
- Novell NDS/eDirectory 8.5 (or higher)

You must also install the shared LDAP library files. Refer to “[LDAP_EXP](#)” on page 435 for installation instructions.

Note

- In Mixed mode, LDAP User Import supports Domain Local Security, Global Security, Domain Local Distribution, Global Distribution, and Universal Distribution groups.
- In Native mode, LDAP User Import supports Domain Local Security, Global Security, Domain Local Distribution, Global Distribution, Universal Security, and Universal Distribution groups.
- If a group contains users from different domains, then only the users of the domain in which the group is created would be imported.

If you use the LDAP User Import tab to govern the importing of users into a library system from Active Directory, you cannot log on to the library via IDM Desktop. And you must set (via IDM Configure) the “Enable LDAP for the whole website” preference for the IDM Web Services web server to “Yes”

If you use the LDAP User Import tab to govern the importing of users into a library system from Novell NDS/eDirectory, you must set (via IDM Configure) the “Use operating system credentials for logon” preference to “Yes” on any machine where you want to log on to the library via IDM Desktop. And you cannot logon to the library via IDM Web Services.

If you use the LDAP User Import tab to govern the importing of users into a library system from an iPlanet/Sun ONE directory, you cannot log on to the library via IDM Desktop. And you must set (via IDM Configure) the “Enable LDAP for the whole website” preference for the Web Services web server to “Yes”.

To configure and start or stop the LDAP User Import tool:

1. Open the Windows Control Panel and start the CS Configuration tool.
2. Select the Installation tab.
3. In the Library System list box, select the name of the library system for which you want to set up LDAP directory synchronization.

4. Select the LDAP 3.0 Server option button as the Directory Server Type. However, if you previously selected another Directory Server Type, stop the Storage Manager, select the Directory Server Type, and then restart the Storage Manager.

Note If you previously selected a different Directory Server Type, a message will pop up to inform you that the Directory Server **group.log** settings will be lost. Click OK to continue or Cancel to retain your earlier Directory Server Type.

5. Select the LDAP User Import tab.

Note If you select the WinNT/Win2000 tab or the XML User Import tab instead of the LDAP User Import tab, a message will pop up to remind you to use the tab corresponding to the Directory Server Type you chose in step 4.

If you want to change your selected Directory Server Type, continue at step 2 of “Using Windows User Import” on page 273 or “Using XML User Import” on page 281, depending on your new selection.

6. In the Interval field, specify how often (in seconds) you want the LDAP Export and CS Import daemons to run. The default is to run hourly (every 3600 seconds). The ideal value for this field will depend on how you plan to use the tool. If you plan to synchronize users and groups and then disable the tool, as recommended by this procedure, you may want to specify a shorter interval (for example, 360 seconds).

If you wish, you can enter an Offset value (in seconds) to offset the daemons' starting time from other processes that run under the Storage Manager. The default offset is 25 minutes (1500 seconds). Refer to “Storage Manager Tab” on page 167 for information about the other Storage Manager processes with configurable offsets.

7. You can specify any valid path and file name for the log this tool produces. The default log file is **Drive:\filenet\LibrarySystem\dirsync.log**.

If the specified file already exists, new logging information will be appended to it. Only changes made since the previous execution are logged.

8. In the Log max size field, enter a maximum size (in kilobytes) for the log file.

9. Specify one of the following:

- the name of the LDAP server and the port (the default is 389) that has the directory containing the groups and users you want to import into your library system.
- the Base Distinguished Name (DN), the Bind DN, and the Bind DN password in the corresponding text boxes. The Base DN points to the directory where the LDAP server entries are located. The Bind DN and Bind DN password are the administrative logon name and password for the LDAP server. For iPlanet, the administrative logon name is usually “cn=Directory Manager”.

10. Check the “Enable user import” box. The daemons will not be enabled until you click Apply or OK.

11. Select the “Select for Organizational Units” check box if you want to use Organizational Units. Clear the check box if you want to use the groupOfUniqueNames or groupOfNames object classes when searching the LDAP directory.

- You must not put a mixture of groups and OUs into the Selected Groups / OU list. The list must contain only groups or only OUs. Otherwise, **dirsync.exe** will detect the selected OUs and use them as the LDAP search filter to import the OUs into the library system. And regardless of the state of the check box, the

selected groups will *not* be imported. The attempt to import them will be recorded as errors in the **dirsync.log** file.

- If you put only groups into the Selected Groups / OUs box, and then select the “Select for Organizational Units” check box, **dirsync.exe** will ignore the status of the check box, and the groups (and their member users) will be imported into the library system.
- If you put only OUs into the Selected Groups / OUs box, and then clear the “Select for Organizational Units” check box, **dirsync.exe** will ignore the status of the check box and import the OUs into the library system.

12. Click Refresh to log in to the LDAP directory server and display the groups (organizational units) in the Available Groups/OUs box.

13. In the Available Groups / OUs box, select the entities you want to import and click the right arrow to move them to the Selected Groups / OUs box. If the Selected box already contains the names of groups / OUs you have previously imported, do not remove them from the Selected box unless you want those groups disabled from the library system.

Notes LDAP User Import does not import groups or their members that are subgroups of LDAP groups. That is, when you specify an LDAP group to be imported/synchronized with the groups and users in a library system, only the LDAP group and its immediate members (users) can be imported, not the subgroups or their members.

LDAP User Import imports LDAP groups and their members that are subgroups of an organizational units. However, organizational units within the OU hierarchy are not imported unless explicitly selected from the CS Configuration tool.

LDAP User Import does not import LDAP users/groups whose names have more than 32 characters (37 characters if the LDAP/ prefix for the user/group name is included).

Caution Subject to the settings of the *Imported Group?* and *Imported User?* properties described in “[User/Group Synchronization Criteria](#)” on page 292, the LDAP User Import tool will disable from the library system any LDAP groups/OUs and users that were imported using this tool and then removed from the Selected Group list for a subsequent execution of this tool.

14. In the Selected Groups / OUs box:

- To control the order in which groups (organizational units) are processed, you can highlight a group (organizational unit) name and use arrow buttons to move the selected name up or down the list.

Note If you have multiple groups (organizational units) selected, the order of processing may be important. If the User Import tool assigns a default library system group (see conditions in “[User/Group Synchronization Criteria](#)” on page 292) it will assign the group created for the first organizational unit processed.

15. Check the “Create preview file only” box to generate a preview transaction file without making any changes to the library system.

You can enter any valid path and file name for the preview transaction file the LDAP User Import tool produces. By default, the preview file is written to **Drive:\filenet\LibrarySystem\preview.log**.

16. Enter a valid pathname for the group-list file the LDAP User Import tool produces. This is the file that contains all the groups that are to be imported into the library system. By default, the group-list file name is **group.log**

17. Once you have all of the options set the way you want them, click Apply. The preview log file will be generated when the daemon is run (as determined by the value you entered in the Run interval field).
18. After the LDAP User Import Tool has executed, review the preview transaction file to verify the changes that will be made to the library system database. Refer to “[User Import Preview Transaction File](#)” on page 307 for a description of the file's content and format.
19. If you are satisfied with the changes outlined in the preview transaction log file, clear the “Create preview file only” check box and click Apply. The next time the daemon runs, your library system users and groups will be synchronized with the selected LDAP directory users and groups.

Using XML User Import

The following is the recommended procedure for using the XML User Import tool. While it is possible to have the XML User Import tool continuously synchronize your library system and LDAP-directory users and groups, before you pursue such a strategy, you should thoroughly review this documentation to understand the implications and risks. You may want to disable LDAP User Import after you have finished importing users and groups.

Before you can import users and groups from an LDAP directory via an XML file, you must set up an LDAP server. The server must be iPlanet Directory Server 5.0 (or higher) running on Sun Solaris 8, HP-UX, Windows 2000 or Windows 2003. Refer to “[Installing an iPlanet LDAP Server](#)” on page 276 for installation information. You must also install the shared LDAP library files. Refer to “[LDAP_EXP](#)” on page 435 for installation instructions.

If you used the XML User Import tab to govern the importing of users into a library system from Active Directory, you cannot log on to the library via IDM Desktop. And you must set (via IDM Configure) the “Enable LDAP for the whole website” preference for the Web Services web server to “Yes”.

If you used the XML User Import tab to govern the importing of users into a library system from Novell NDS/eDirectory, you cannot log on to the library via IDM Desktop. And you must set (via IDM Configure) the “Enable LDAP for the whole website” preference for the Web Services web server to “Yes”.

If you used the XML User Import tab to configure the settings that govern the importing of users into a library system from an iPlanet/Sun ONE directory, you cannot log on to the library via IDM Desktop. And you must set (via IDM Configure) the “Enable LDAP for the whole website” preference for the Web Services web server to “Yes”.

To configure and start or stop the XML User Import tool:

1. Open the Windows Control Panel and start the CS Configuration tool.
2. Select the Installation tab.
3. In the Library System list box, select the name of the library system for which you want to set up LDAP directory synchronization.
4. Select the FileNet XML File option button as the Directory Server Type. However, if you previously selected another Directory Server Type, stop the Storage Manager, select the Directory Server Type, and then restart the Storage Manager.

Note If you previously selected a different Directory Server Type, a message will pop up to inform you that the Directory Server **group.log** settings will be lost. Click OK to continue or Cancel to retain your earlier Directory Server Type.

5. Select the XML User Import tab.

Note If you select the WinNT/Win2000 tab or the XML User Import tab instead of the LDAP User Import tab, a message will pop up to remind you to use the tab corresponding to the Directory Server Type you chose in step 4.

If you want to change your selected Directory Server Type, continue at step 2 of “[Using Windows User Import](#)” on page 273 or “[Using LDAP User Import](#)” on page 277, depending on your new selection.

6. In the Interval field, specify how often (in seconds) you want the LDAP Export and CS Import daemons to run. The default is to run hourly (every 3600 seconds). The ideal value for this field will depend on how you plan to use the tool. If you plan to synchronize users and groups and then disable the tool, as recommended by this procedure, you may want to specify a shorter interval (for example, 360 seconds).

If you wish, you can enter an Offset value (in seconds) to offset the daemons' starting time from other processes that run under the Storage Manager. The default offset is 25 minutes (1500 seconds). Refer to [“Storage Manager Tab” on page 167](#) for information about the other Storage Manager processes with configurable offsets.

7. You can specify any valid path and file name for the log this tool produces. The default log file is **Drive:\filenet\LibrarySystem\dirsync.log**.

If the specified file already exists, new logging information will be appended to it. Only changes made since the previous execution are logged.

8. In the Log max size field, enter a maximum size (in kilobytes) for the log file.
9. Specify the absolute pathname of an XML file generated by the LDAP_EXP utility (refer to [“LDAP_EXP” on page 435](#)). The XML file must *not* be on a mapped drive.
10. Check the “Enable user import” box. The daemons will not be enabled until you click Apply or OK.
11. Click Refresh to log onto the LDAP directory server and display the groups in the Available XML Groups box.
12. In the Available XML Groups box, select the groups you want to import and click the right arrow to move them to the Selected XML Groups box. If the Selected box already contains the names of groups you have previously imported, do not remove them from the Selected box unless you want those groups disabled from the library system.

Notes XML User Import does not import LDAP subgroups or their members. That is, when you specify an LDAP group to be imported/synchronized with the groups and users in a library system, only the LDAP group and its immediate members (users) can be imported, not the subgroups or their members.

XML User Import does not import LDAP users/groups whose names have more than 32 characters (37 characters if the LDAP/ prefix for the user/group name is included).

Caution Subject to the settings of the *Imported Group?* and *Imported User?* properties described in [“User/Group Synchronization Criteria” on page 292](#), the XML User Import tool will disable from the library system any LDAP groups and users that were imported using this tool and then removed from the Selected XML Group list for a subsequent execution of this tool.

Also, any user/group previously imported from a Windows NT domain or Windows 2000 or Windows 2003 Active Directory will be disabled by the Windows User Import daemon **dirsync** if its Imported User?/Imported Group? property value is *NTLM*.

13. In the Selected XML Groups box:
 - a. To control the order in which groups are processed, you can highlight a group name and use arrow buttons to move the selected name up or down the list.

Note If you have multiple groups selected, the order of processing may be important. If the User Import tool assigns a default library system group (see conditions in [“User/Group Synchronization Criteria” on page 292](#)) it will assign the first processed.

14. Check the “Create preview file only” box to generate a preview transaction file without making any changes to the library system.

You can enter any valid path and file name for the preview transaction file the XML User Import tool produces. By default, the preview file is written to *Drive:\filenet\LibrarySystem\preview.log*.

15. Enter a valid pathname for the group-list file the XML User Import tool produces. This is the file that contains all the groups that are to be imported into the library system. By default, the group-list file name is **group.log**
16. Once you have all of the options set the way you want them, click Apply. The preview log file will be generated when the daemon is run (as determined by the value you entered in the Run interval field).
17. After the XML User Import Tool has executed, review the preview transaction file to verify the changes that will be made to the library system database. Refer to “[User Import Preview Transaction File](#)” on page 307 for a description of the file's content and format.
18. If you are satisfied with the changes outlined in the preview transaction log file, clear the “Create preview file only” check box and click Apply. The next time the daemon runs, your library system users and groups will be synchronized with the selected LDAP directory users and groups.

Passwords for Imported Users

The User Import tools assign each imported user an encrypted, unknown password. The imported users can then log on to the library system using the Network Logon option. Refer to [“Using Network Logons” on page 195](#) for information about this feature.

If your site does not use the Network Logon option, the password for each imported user must be changed immediately after import. As system administrator, you can do this manually by changing each User object's Password property, or you may instead want to run a Bulk Modify process to set the Password property to a known value or blank, and then notify all imported users to change their password to a unique value only they know. Refer to *Content Services Admin Tools Online Help* for procedures.

User Import Usage Guidelines

The usage guidelines listed here are divided into four sections, according to the method you use to import users and groups. The guidelines listed in each of the first three sections are unique to that method. The guidelines in the fourth section are common to all the User Import methods.

- “Windows User Import” on page 285
- “LDAP User Import” on page 286
- “XML User Import” on page 287
- “Windows/LDAP/XML User Import” on page 289

References to “User Import” in the guidelines listed in “Imported GroupType? and Imported User? Properties” on page 290 apply to Windows User Import, LDAP User Import, and XML User Import.

Windows User Import

- Windows User Import consists of a single program, **dirsync.exe**, which runs on the initial storage server as a daemon under a selected library system's Storage Manager. To function correctly, this file must be located in the *Drive:\filenet\LibrarySystem\pbin* directory.
- Windows User Import supports Windows mixed-mode domains, but not native-mode domains.
- The library system whose users and groups you want to synchronize with those in a Windows NT domain must be attached to that domain.
- To synchronize users and groups in a library system with a Windows 2000 or Windows 2003 Active Directory, you should be logged onto the domain as a user with at least Print Operator privileges. Your username should also be in the Administrators group on the CS server where the library system is installed.
- Windows NT groups belonging to a domain that does not have a trust relationship with the domain of your initial storage server are not available for selection via Windows User Import. In other words, if the server you are configuring is in a different domain than the one you select in the Domain selection box, and if your server's domain does not have a trust relationship with the selected domain, you will not be able to see that domain's groups in the Available Groups list.
- To enable the tool to import groups from “non-trusting” domains you can establish a two-way trust between the domains, or you can create a temporary trust by mapping a drive to the other domain's Primary Domain Controller share device.
- If the same group name exists in multiple NT domains, Windows User Import will treat the domain groups as the same group when determining a match with your library system group names. In other words, if a domain group name exists in multiple domains, and you select those domain groups for processing, Windows User Import will treat them as a single group. This also holds true for userids that exist in multiple domains; they are treated as the same userid.
- If you attempt to rename a user or a group in a domain while the Windows User Import daemon is running, the user with the former name will be disabled and a new user with the new name will be created. To avoid introducing such inconsistencies between the users/groups in the domain and those in Content Services, use this procedure:

To rename a user or group when using Windows User Import:

- a. Shut down the Windows User Import daemon.
- b. Use your network tools to rename the affected user or group in the domain. Be sure to rename *all* applicable fields. For groups, the applicable fields are the Group Name and the Pre-Windows 2000 or Windows 2003 Group Name. For users, the applicable fields are the User Name, Full Name, Display Name, User Logon Name, and Pre-Windows 2000 or Windows 2003 User Name.
- c. Manually rename (via CS Admin Tools) the corresponding User or Group objects in Content Services.

Be sure to rename *all* applicable property values in the User and Group objects maintained by CS Admin Tools. For Group objects, the applicable properties are the Group Name and Display Name. For User objects, the applicable properties are the User Name and the Display Name.

- d. Restart the Windows User Import daemon.

If you renamed the domain user or group without first shutting down the User Import daemon, you can restore consistency by performing steps a, b, and d.

- Windows User Import ignores domain userids that end with the dollar sign (\$) character. It also ignores userids that begin with:
 - IUSR_
 - VUSR_
 - SMS
 - SQL
 - IWAN_
- The default location for the Windows User Import log file is *Drive:\filenet\LibrarySystem\dirsync.log*. You can specify a different directory or file name if you wish, but you should exercise caution if more than one library system is installed on the initial storage server. Do not direct User Import logs for multiple library systems to a single location.
- When you add a user to a network domain, then add that user to a group, and then delete the user from the NT domain, the userid is *not* automatically deleted from the groups. You must manually delete the user from the groups.

LDAP User Import

- LDAP User Import consists of a single program, **dirsync.exe**, that runs on the initial storage server as a daemon under a selected library system's Storage Manager. This program is located in the *Drive:\filenet\LibrarySystem\pbin* directory.
- To import users and groups/OUs from an LDAP directory, specify LDAP directory information in the LDAP User Import tab of the CS Configuration tool to obtain a list of groups or OUs to import. See “LDAP User Import Tab” on page 180
- LDAP users and groups/OUs removed from the LDAP directory will be subsequently disabled in Content Services.

- An LDAP user belonging to more than one group/OU will have as its Default Group the group/OU appearing closest to the beginning of the list in the Selected Groups/OUs box in the LDAP User Import tab of the CS Configuration tool.
- LDAP groups/OUs removed from the CS Selected Groups/OUs box in the LDAP User Import tab of the CS Configuration tool will be subsequently disabled in Content Services.
- If you attempt to rename a user or a group in an LDAP directory while the LDAP User Import daemon is running, the daemon will not be able to automatically rename the corresponding user or group object in the Content Services library system. Instead, the user/group with the former name will be disabled and a user/group with the new name will be created.

To rename a user or group (or organizational unit) when using LDAP User Import:

- Disable the LDAP User Import daemon.
- Use your LDAP administrative tools to rename the affected user or group in the LDAP directory.

Be sure to rename *all* applicable fields. For groups and organizational units, the applicable field is Name. For users, the applicable fields are the Name, Surname, and UniqueID.
- Manually rename (via CS Admin Tools) the corresponding User or Group objects in Content Services.

Be sure to rename *all* applicable property values in the User and Group objects maintained by CS Admin Tools. For Group objects, the applicable properties are the Group Name and Display Name. For User objects, the applicable properties are the User Name and the Display Name.
- Restart the LDAP User Import daemon.

If you renamed the LDAP user or group without first disabling the LDAP User Import daemon, you can restore consistency by performing steps a, b, and d.

- The default location for the LDAP User Import log file is *Drive:\filenet\LibrarySystem\dirsync.log*. You can specify a different directory or file name if you wish, but you should exercise caution if more than one library system is installed on the initial storage server. Do not direct User Import logs for multiple library systems to a single location.
- If you add a user to an LDAP directory, and then add that user to a group, and finally delete the user from the LDAP directory, the user is *not* automatically deleted from the groups. You must manually delete the users from the groups.
- When you add a user to an LDAP directory, first add that user to a group, and then delete the user from the NT domain.

XML User Import

- XML User Import consists of two standalone programs, **ldap_exp** and **dirsync.exe**, which run on the initial storage server under a selected library system's Storage Manager. **Ldap_exp** queries an LDAP directory for a list of users, groups, and group memberships, which it writes to an XML file that is then used by **dirsync.exe** to import the LDAP users and groups into a Content Services library system.
- **Ldap_exp** can be invoked from a Windows DOS command line, or a script/batch file operation. **Ldap_exp** can also be scheduled via the Schedule Service (on Windows NT), or Task Scheduler (on Windows 2000 or Windows 2003).

- The **dirsync.exe** program can be invoked from a Windows DOS command line, a script/batch file operation, or automatically at specified intervals (set via the XML User Import tab in the CS Configuration tool) if the Content Services Storage Manager is running.
- The XML file processed by **dirsync.exe** must *not* reside on a mapped drive.
- All the required parameters for **ldap_exp** are supplied via the command-line argument list. See “LDAP_EXP” on page 435 for details.
- XML User Import imports users and groups only from iPlanet Directory Server 5.0 (or higher), running on a Sun Solaris 8.
- XML User Import does not import LDAP subgroups or their members. That is, when you specify an LDAP group to be imported/synchronized with the groups and users in a library system, only the LDAP group and its immediate members (users) can be imported, not the subgroups or their members.
- XML User Import does not import LDAP users/groups whose names have more than 32 characters (37 characters if the LDAP/ prefix for the user/group name is included).
- You should not schedule **dirsync.exe** to run until after **ldap_exp** has run to completion. (The running time of **ldap_exp** depends on the number of users it is exporting.) Otherwise, **dirsync.exe** will be reading the **ldap_exp.xml** file before LDAP_EXP has finished generating it.
- The specific attribute in an XML file that is used to authenticate library system logons with a directory server depends on the directory server from which **ldap_exp** obtains user/group information:
 - If the directory server is iPlanet, the attribute is **uid**.
 - if the directory server is Novell eDirectory (NDS), the attribute is **cn**.
 - If the directory server is Active Directory, the attribute is **SAMAccountName**.
- Do not attempt to rename users or groups on a server where a utility is generating an XML file of users and groups without first disabling the XML User Import daemon. Otherwise, the daemon will not be able to automatically rename the corresponding user or group object in the Content Services library. Instead, the user/group with the former name will be disabled and a user/group with the new name will be created.

To rename a user or group imported via an XML file:

- a. Disable the XML User Import daemon by clearing the “Enable user import” checkbox in the XML User Import tab of the CS Configuration tool.
- b. Rerun the utility on the LDAP server that generates the XML file (which now contains the renamed users and groups). For groups, the applicable field is the Group Name. For users, the applicable fields are the First Name, Last Name, Common Name(s), and User ID.
- c. Manually rename (via CS Admin Tools) the corresponding User or Group objects in Content Services.

Be sure to rename *all* applicable property values in the User and Group objects maintained by CS Admin Tools. For Group objects, the applicable properties are the Group Name and Display Name. For User objects, the applicable properties are the User Name and the Display Name.
- d. Re-enable the XML User Import daemon by selecting the “Enable user import” checkbox.

- When you are importing many users and groups from an LDAP directory into a Content Services library system, the size of the XML file that LDAP_EXP generates may be several megabytes. And if LDAP_EXP runs with the -v option (verbose log output), then the log file **ldap_exp.log** generated by LDAP_EXP can also be large (several hundred kilobytes to a megabyte or more).

Since **dirsync.exe** renames the XML file each time it runs, and appends entries to its log file, you will accumulate many XML files and a lengthy log file (if log output is verbose) over time if LDAP_EXP and CS_IMPORT run periodically. Thus, if these tools run on a Windows machine, you should make sure that the XML files are created in a partition with sufficient size.

XML User Import using LDIF (LDAP Data Interchange Format) Format Files

LDIF (LDAP Data Interchange Format) is a file format based on Internet Engineering Task Force (IETF) Request for Comments (RFC) 2849 for importing and exporting data from LDAP directories.

Every LDAP system has utilities to create LDIF files. These files can be created in batches. FileNet Image Services provides a tool **fn_ldif_xfer** that converts the LDIF format files into FileNet proprietary XML format file. After this XML file is created, CS_IMPORT can be used to import users/groups to CS library using the batch mode flag. In batch mode, only add (or insert) operation is performed with no delete operation, i.e., the users and groups are added to the library system provided they do not exist in the system. If the users and groups already exist, then they are updated. The groups which do not have any members are added as disabled groups.

Windows/LDAP/XML User Import

- User Import is designed to run until it has, in a single pass, completely synchronized your library system group and user information with the selected domain (LDAP Directory) group and user information. With each successive execution, User Import keeps the library system group and user information synchronized with the source Windows domain or LDAP directory information.
- Before performing administrative tasks with the Paragon Content Services Admin Tools, you should verify (via the CS Configuration tool) that User Import is disabled. You can then perform any necessary bulk add or administrative tasks, and then use the administrative tools to verify that your library system looks okay.
- Before you enable User Import to perform actual processing on your library-system database, we highly recommend that you generate a Preview Transaction File to see what changes the daemon will make. If the Preview Transaction File reports errors or changes that you do not want to make, you should resolve them before enabling User Import for actual production processing. Additionally, the log files (**dirsync.log** for Windows/LDAP/XML User Import) should be reviewed for potential issues, which must be resolved prior to running User Import.
- User Import's default execution interval is to run every hour. Be sure to set this interval to a value appropriate for your network. For example, you may want to set it to 86,400 seconds to run once a day.
- After making any changes to settings in the WinNT/Win2000, LDAP, or XML User Import tab, you must stop and restart the Storage Manager.
- If you are also using replication services on the library system, the following order is recommended:
 - For new library system installations, run User Import before enabling replication on the library system.
 - When upgrading a library system that did not participate in any user/group synchronization, once the upgrade is complete you must rename all library system userids to match the userids and group names in

the NT domain or LDAP directory. Next, enable replication before setting the User Import? property to “NTLM”, “LDAP”, or “XML”, and activating the User Import tool.

- User Import does not synchronize passwords between the library system and the Windows domain or LDAP directory. Any attempt to do so would be a security breach.
- If your library system has an existing user named Admin, that user does *not* become a member of newly imported Windows domain or LDAP groups that also contain a user named Admin.
- If you import groups and users via one directory synchronization method (LDAP, for example), and then change to a another method (Active Directory, for example), the users and groups imported by the first method will be ignored by the new method. That is, the users and groups will not be disabled, even though they do not appear in the new directory.

If you want to disable a set of users or groups, you need to return to the previous directory synchronization method, delete the users and groups from their directory, and allow synchronization to restore consistency with the users and groups in the library system. Then you can change to the new directory synchronization method.

- If **dirsync** imports a group and its member users from a directory server into a library system, and a user's membership in the group on the directory server is subsequently removed, **dirsync** is supposed to remove that membership in the imported CS group. If the imported user is a member of exactly one imported group in the library, then **dirsync** works as it should.

If, however, the imported user is a member of more than one imported group then, because of a bug (see FNDTS00071097), the imported user remains in the imported group after the user's membership in the group on the directory server has been removed. In this case, the library system administrator must manually remove the CS membership from the imported group.

Imported Group? and Imported User? Properties

- A primary purpose of the Imported Group? and Imported User? properties is to support “single-pass” types of usage cases for User Import. For example, in such a scenario you might:
 - a. Run a single pass of User Import to import a particular network domain group and its users (administrators, for example) into your library system.
 - b. Shut down the User Import Tool.
 - c. Set the Imported Group? and Imported User? properties to No for the administrators group and its users, to specify that you will perform ongoing maintenance of them manually.
 - d. Enable User Import to import other groups/users and allow the daemon to continue running to maintain them automatically.
- Exercise caution when changing the value of the Imported Group? and Imported User? properties. For example:
 - For a library-system group or user that was not added via the User Import Tool, if you set the Imported Group? or Imported User? property to “XML”, “NTLM”, or “LDAP” and the group or user in question *does* exist in the network domain, the next pass of the User Import daemon will overwrite its existing library system information with information from the Group Defaults, unless the property is set to “XML NOUPDATE”, “NTLM NOUPDATE”, or “LDAP NOUPDATE”. See [“Imported Group Defaults” on page 184](#).

- For a library-system group or user that was not added via the User Import Tool, if you set the Imported Group? or Imported User? property value to anything other than “No” and the group or user in question does not exist in the network domain or LDAP directory, the next pass of the User Import daemon will disable the user/group, preventing it from logging onto the library system. Also, this property is not automatically set to “Yes” for the Administrator user. It is left blank.
- While a Group or User object's Imported Group? or Imported User? property is set to anything other than “No”, you should not delete or rename the group or user manually via the CS Admin Tools. You should first set the property to “No”, then make the change. You can then, if you wish, set this property back to its former value to allow the User Import Tool to perform ongoing synchronization of the group or user.
- Prior to running the User Import tool, IDM Desktop must be configured to enable the “Use operating system credentials for logon” property only for Novell eDirectory LDAP directories or Microsoft Active Directory.
- It is recommended that you not enable the Import User? property for any Administrator users.

User/Group Synchronization Criteria

The synchronization criteria described in this section apply to groups based on either Windows network domain groups or on LDAP organizational units. References to “User Import” apply to either Windows Domain User Import or LDAP User Import.

For each domain/LDAP group the User Import Tool first determines whether the group name already exists in the library system database.

Note The following changes to library system Group and User objects are subject to the setting of the Imported Group? and Imported User? properties (described below). These properties are described in the next topic.

- If the domain/LDAP group name does *not* already exist in the library system, User Import adds it and processes its userid list as follows:
 - For a userid that is new to the library system, the domain/LDAP group being processed is assigned as the user's default group in the library system.
 - For a userid that already exists in the library system, the name of the domain/LDAP group/OU being processed is added to the user's group membership list in the library system if the Imported User? property value is anything other than “No” and if the user belongs to the domain/LDAP group/OU being processed.
- If the group name *does* already exist in the library system, User Import merges the domain/LDAP group's user list with the library system group's users list, processing userids as follows:
 - For userids that exist in the domain/LDAP group but not in the library system, the users are added to the library system and the domain/LDAP group being processed is assigned as the default library system group.
 - For userids that exist in the domain/LDAP group and in the library system, their associated library system group membership list is updated.
 - For userids that exist in the library system group but not in the domain/LDAP group, the user is removed as a member from the library system group.
 - For userids that exist in the library system group but not in the domain/LDAP group or in the domain/LDAP userid list, and the User object's Imported User? property value is anything other than “No”, the user is disabled from logging onto the library system.

For details about what group and user information the User Import Tool updates, and where it gets the update information, refer to “[User Import Preview Transaction File](#)” on page 307.

The Imported Group? and Imported User? Properties

Group and User objects each have a property that determines whether User Import will keep them synchronized with their respective network domain sources or LDAP organizational units.

- The Group object **Imported Group?** property indicates whether the group's name and property information are imported from, and regularly synchronized with, the network domain group information. This property is initially set to one of the following if the group is imported by the User Import tool:
 - “XML” or “XML NOUPDATE” if an XML file (generated by **ldap_exp**) was used for retrieving a list of LDAP OUs or groups.

- “NTLM” or “NTLM NOUPDATE” if Active Directory domains are specified to import groups from.
- “LDAP” or “LDAP NOUPDATE” if LDAP Server, BaseDN, BindDN, Bind Password are specified for retrieving a list of OUs or groups.

The Imported Group? property value is set to “No” if the group is added and maintained manually via CS Admin Tools. Refer to *CS Explorer Online Help* for information about how you can change the value manually.

- The User object **Imported User?** property indicates whether the user's name and property information are imported from, and regularly synchronized with, the network domain user (LDAP organizational unit) information. This property is initially set to one of the following if the user is imported by the User Import tool:
 - “XML” or “XML NOUPDATE” if an XML file (generated from **ldap_exp**) was used for retrieving a list of LDAP OUs or groups
 - “NTLM” or “NTLM NOUPDATE” if Active Directory domains are specified to import groups from
 - “LDAP” or “LDAP NOUPDATE” if LDAP Server, BaseDN, BindDN, Bind Password are specified for retrieving a list of OUs or groups.

The Imported User? property value is set to “No” if the group is added and maintained manually via CS Admin Tools. Refer to *CS Explorer Online Help* for information about how you can change the value manually.

Note FNDTS00070376: If you change the password (via CS Admin Tools) of an imported user (group) from its imported encrypted value to a plaintext value, and the Imported User Type (Imported Group Type) property value of the imported user (group) is set to XML, LDAP, or NTLM, **dirsync** will fail to maintain the synchronization of the imported-user (group) password with the password of the user (group) maintained on the network directory server. This affects user logons to library systems as follows:

- Users imported via LDAP User Import or XML User Import can no longer authenticate to their LDAP directory. However, they can still log on with the conventional, non-LDAP-enabled login dialog in Web Services.
- Users imported via Window NT/Windows 2000 or Windows 2003 can still log on using unified logons.

For guidelines on using these properties, refer to “[User Import Usage Guidelines](#)” on page 285. For instructions on how to change the value of these properties, refer to the *CS Explorer Online Help*.

Windows User Import Usage Example

The following is a simple usage scenario showing how a site might import certain groups and users from two network domains (SAMPLECO1 and SAMPLECO2) into a newly installed library system, and then keep the library system up to date as changes are made to the domain group and user information.

The composition of the two domains is:

| SAMPLECO1 Domain | SAMPLECO2 Domain |
|---|--|
| FINANCE Group | FINANCE Group |
| <ul style="list-style-type: none"> • LAtbrough • JSmithers • LBose | <ul style="list-style-type: none"> • CJohnson • GMathers • DJames |
| MANAGERS Group | MANAGERS Group |
| <ul style="list-style-type: none"> • CWhithers • LAtbrough | <ul style="list-style-type: none"> • HALLman • CJohnson |

The administrator installs a library system named FINANCE on server FINANCE1 in domain SAMPLECO1. The intent is to grant the members of the FINANCE and MANAGERS groups in both domains access to the FINANCE library system. To achieve this, the administrator sets up the User Import Tool for library system FINANCE, and enters SAMPLECO1 in the Domain: text box. The domain groups, including FINANCE and MANAGERS are displayed in the Available Groups list box. These two groups are highlighted and moved into the Selected Groups list box by clicking the arrow button. Repeat for the SAMPLECO2 domain.

- SAMPLECO1\FINANCE
- SAMPLECO1\MANAGERS
- SAMPLECO2\FINANCE
- SAMPLECO2\MANAGERS

The first execution of the daemon creates two new library system groups:

```
FINANCE, Sampleco's two finance groups
MANAGERS, Sampleco's two manager groups
```

Also in the first execution, eight new library system users are created:

```
LAtbrough, Len Atbrough, Finance Manager, SAMPLECO1\LAtbrough, FINANCE, MANAGERS
JSmithers, Julie Smithers, Company Finance Asst., SAMPLECO1\JSmithers, FINANCE
LBose, Linda Bose, SAMPLECO1 Finance Asst., SAMPLECO1\LBose, FINANCE
CJohnson, Chris Johnson, Finance Manager, SAMPLECO2\CJohnson, FINANCE, MANAGERS
GMathers, Greg Mathers, Finance Asst., SAMPLECO2\GMathers, FINANCE
DJames, Don James, Finance Asst., SAMPLECO2\Djames, FINANCE
CWhithers, Craig Whithers, SAMPLECO1 Boss, SAMPLECO1\CWhithers, MANAGERS
CJohnson, Chris Johnson, SAMPLECO2 Boss, SAMPLECO2\CJohnson, MANAGERS
```

The system administrator leaves the User Import Tool enabled, configured to maintain the FINANCE library system's group and user information.

Now, there's a minor reorganization in the SAMPLECO1 domain. Len Atbrough has quit and Craig Whithers has temporarily taken over as finance manager. The administrator makes the necessary changes in the network domain and User Import then processes the new information and updates the library system:

- User Latbrough is removed from the library system
- User CWhithers is added to the library system's FINANCE group

The next change in the SAMPLECO1 domain comes as a result of Charles Henry being hired as the new finance manager. Craig Whithers now wants out of the FINANCE group. The administrator makes the necessary changes to the domain, and User Import then updates the library system:

- User CHenry is added to the library system and made a member of the FINANCE group and the MANAGERS group
- User CWhithers is removed from the library system FINANCE group

Windows User Import Log File Example

The log below illustrates how a new library system might look after a first-time run of the Windows User Import Tool configured to process three domain groups:

- SRVDEVTEST\Domain Users
- SRVDEVTEST\TestGroup
- SRVDEVTEST\TestGroup2

```

15:43|30/Jun/1999|=====
15:43|30/Jun/1999|Session established on server LKLINE3, system ENG1.
15:43|30/Jun/1999|Adding Group: Domain Users, All domain users.
15:43|30/Jun/1999|New User Administrator added for CS Group Domain Users.
15:43|30/Jun/1999|New User bjones added for CS Group Domain Users.
15:43|30/Jun/1999|New User bsmith added for CS Group Domain Users.
15:43|30/Jun/1999|New User cthomas added for CS Group Domain Users.
15:43|30/Jun/1999|New User dgates added for CS Group Domain Users.
15:43|30/Jun/1999|New User jallen added for CS Group Domain Users.
15:43|30/Jun/1999|New User jhancock added for CS Group Domain Users.
15:43|30/Jun/1999|New User mtodd added for CS Group Domain Users.
15:43|30/Jun/1999|New User rrandel added for CS Group Domain Users.
15:43|30/Jun/1999|New User tcunningham added for CS Group Domain Users.
15:43|30/Jun/1999|Adding Group: TestGroup, Lyle's Test Group.
15:43|30/Jun/1999|New User bjones added to CS Group TestGroup.
15:43|30/Jun/1999|New User bsmith added to CS Group TestGroup.
15:43|30/Jun/1999|New User cthomas added to CS Group TestGroup.
15:43|30/Jun/1999|New User jallen added to CS Group TestGroup.
15:43|30/Jun/1999|New User rrandel added to CS Group TestGroup.
15:43|30/Jun/1999|New User tcunningham added to CS Group TestGroup.
15:43|30/Jun/1999|Adding Group: TestGroup2, Lyle's Test Group #2 (GLOBAL).
15:43|30/Jun/1999|New User dgates added to CS Group TestGroup2.
15:43|30/Jun/1999|New User rrandel added to CS Group TestGroup2.
15:43|30/Jun/1999|New User tcunningham added to CS Group TestGroup2.
15:43|30/Jun/1999|Session disconnected on server LKLINE3, system ENG1.
15:43|30/Jun/1999|=====
  
```


Sample LDAP Export XML File

The following XML file is a sample of the output generated by the LDAP_EXP utility. Refer to “LDAP_EXP” on page 435 and “LDAP Export XML Schema” on page 300 for more information about the LDAP_EXP utility and the schema that represents the utility's XML output.

```
<?xml version="1.0" encoding="UTF-8"?>
<principals xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="filenet_ulf_v1.xsd">
<users>
  <user>
    <dn>uid=JDover,ou=People, dc=idmds,dc=com</dn>
    <id>LDAP/JDover</id>
    <cn>Jill Dover</cn>
  </user>
  <user>
    <dn>uid=JJohnson,ou=People, dc=idmds,dc=com</dn>
    <id>LDAP/JJohnson</id>
    <cn>Jack Johnson</cn>
  </user>
  <user>
    <dn>uid=JJones,ou=People, dc=idmds,dc=com</dn>
    <id>LDAP/JJones</id>
    <cn>John Jones</cn>
  </user>
  <user>
    <dn>uid=JSmith,ou=People, dc=idmds,dc=com</dn>
    <id>LDAP/JSmith</id>
    <cn>James Smith</cn>
  </user>
</users>
<groups>
  <group>
    <dn>cn=Group 1,ou=Groups, dc=idmds,dc=com</dn>
    <id>LDAP/Group 1</id>
    <cn>Group 1</cn>
    <description>Group named "Group 1"</description>
  </group>
  <group>
    <dn>cn=Group c,ou=Groups, dc=idmds,dc=com</dn>
    <id>LDAP/Group c</id>
    <cn>Group c</cn>
    <description>Group named "Group c"</description>
  </group>
  <group>
    <dn>cn=Group b,ou=Groups, dc=idmds,dc=com</dn>
    <id>LDAP/Group b</id>
    <cn>Group b</cn>
    <description>Group named "Group b"</description>
  </group>
  <group>
    <dn>cn=Group a,ou=Groups, dc=idmds,dc=com</dn>
    <id>LDAP/Group a</id>
```

```

    <cn>Group a</cn>
    <description>Group named "Group a"</description>
</group>
<group>
    <dn>cn=Group i,ou=People, dc=idmds,dc=com</dn>
    <id>LDAP/Group i</id>
    <cn>Group i</cn>
    <description>Group named "Group i"</description>
</group>
</groups>
<members>
    <member>
        <groupName>cn=Group 1,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>cn=Group b,ou=Groups, dc=idmds,dc=com</memberName>
        <memberType>group</memberType>
    </member>
    <member>
        <groupName>cn=Group 1,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>uid=JDover,ou=People, dc=idmds,dc=com</memberName>
        <memberType>user</memberType>
    </member>
    <member>
        <groupName>cn=Group 1,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>uid=JJones,ou=People, dc=idmds,dc=com</memberName>
        <memberType>user</memberType>
    </member>
    <member>
        <groupName>cn=Group 1,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>uid=JJJohnson,ou=People, dc=idmds,dc=com</memberName>
        <memberType>user</memberType>
    </member>
    <member>
        <groupName>cn=Group c,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>cn=Group i,ou=People, dc=idmds,dc=com</memberName>
        <memberType>group</memberType>
    </member>
    <member>
        <groupName>cn=Group c,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>uid=JJJohnson,ou=People, dc=idmds,dc=com</memberName>
        <memberType>user</memberType>
    </member>
    <member>
        <groupName>cn=Group b,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>uid=JJones,ou=People, dc=idmds,dc=com</memberName>
        <memberType>user</memberType>
    </member>
    <member>
        <groupName>cn=Group a,ou=Groups, dc=idmds,dc=com</groupName>
        <memberName>cn=Group i,ou=People, dc=idmds,dc=com</memberName>
        <memberType>group</memberType>
    </member>
    <member>
        <groupName>cn=Group a,ou=Groups, dc=idmds,dc=com</groupName>

```

```
<memberName>uid=JJohnson,ou=People, dc=idmids,dc=com</memberName>
<memberType>user</memberType>
</member>
<member>
  <groupName>cn=Group a,ou=Groups, dc=idmids,dc=com</groupName>
  <memberName>uid=JJones,ou=People, dc=idmids,dc=com</memberName>
  <memberType>user</memberType>
</member>
<member>
  <groupName>cn=Group a,ou=Groups, dc=idmids,dc=com</groupName>
  <memberName>uid=JSmith,ou=People, dc=idmids,dc=com</memberName>
  <memberType>user</memberType>
</member>
<member>
  <groupName>cn=Group i,ou=People, dc=idmids,dc=com</groupName>
  <memberName>uid=JSmith,ou=People, dc=idmids,dc=com</memberName>
  <memberType>user</memberType>
</member>
</members>
</principals>
```

LDAP Export XML Schema

This topic describes the XML schema for the LDAP export data files that are generated by the LDAP_EXP utility.

The topic includes the following subtopics:

- “Conventions” on page 300
- “Element List” on page 300
- “Element Descriptions” on page 301
- “Schema Source” on page 304

Conventions

This document uses the following conventions:

| Occurrence Indicators | |
|-----------------------|---|
| <i>none</i> | Required and non-repeatable. |
| ? | Optional and non-repeatable. |
| * | Optional and repeatable. |
| + | Required and repeatable. |
| , | Sequence. |
| | Choice. |
| : | Identifies a default after a list of enumerated values. |
| (...) | Group. |

Element List

This schema includes the following elements:

| | | |
|-------------|------------|------------|
| cn | groups | members |
| dn | id | memberType |
| description | identity | principals |
| group | member | user |
| groupName | memberName | users |

Element Descriptions

| | |
|-------------|--|
| cn | |
| description | LDAP common name (CN). |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | identity (extension to the user and group elements) |
| content | string |

| | |
|-------------|--|
| dn | |
| description | LDAP distinguished name (DN). |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | identity (extension to the user and group elements) |
| content | string |

| | |
|--------------------|--|
| description | |
| description | LDAP description. |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | identity (extension to the user and group elements) |
| content | string |

| | |
|--------------|--|
| group | |
| description | A collection of elements that identify a group. |
| attributes | <i>none</i> |
| uses | dn, id, cn, description? (as derived by the identity extension) |
| used by | groups |
| content | sequence |

| | |
|------------------|---|
| groupName | |
| description | LDAP distinguished name (DN) for a group. |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | member |
| content | string |

| | |
|---------------|-------------------------------|
| groups | |
| description | Collection of group elements. |
| attributes | <i>none</i> |
| uses | group* |
| used by | principals |
| content | sequence |

| | |
|-------------|--|
| id | |
| description | Unique ID. Typically, an LDAP uid or cn attribute value. |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | identity (extension to the user and group elements) |
| content | string |

| | |
|-----------------|---|
| identity | |
| description | Collection of elements that describe a user or group. |
| attributes | <i>none</i> |
| uses | dn, id, cn, description? |
| used by | <i>none</i> (extends user and group) |
| content | sequence |

| | |
|---------------|--|
| member | |
| description | Collection of elements that describe a member. A member can be either a user or group. |
| attributes | <i>none</i> |
| uses | groupName, memberName, memberType |
| used by | members |
| content | sequence |

| | |
|-------------------|--|
| memberName | |
| description | LDAP distinguished name (DN) for a member. |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | member |
| content | string |

| | |
|----------------|--------------------------------|
| members | |
| description | Collection of member elements. |
| attributes | <i>none</i> |
| uses | member* |
| used by | principals |
| content | sequence |

| | |
|-------------------|--|
| memberType | |
| description | Specifies whether a member is a user or group. |
| attributes | <i>none</i> |
| uses | <i>none</i> |
| used by | member |
| content | string (Restricted to: group user) |

| | |
|-------------------|--|
| principals | |
| description | Root element of an LDAP export XML file. |
| attributes | <i>none</i> |
| uses | users, groups, members |
| used by | <i>none</i> |
| content | sequence |

| | |
|-------------|--|
| user | |
| description | A collection of elements that identify a user. |
| attributes | <i>none</i> |
| uses | dn, id, cn, description? (as derived by the identity extension) |
| used by | users |
| content | sequence |

| | |
|--------------|------------------------------|
| users | |
| description | Collection of user elements. |
| attributes | <i>none</i> |
| uses | user* |
| used by | principals |
| content | sequence |

Schema Source

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
elementFormDefault="qualified">
  <xs:element name="principals">
    <xs:complexType>
      <xs:sequence>
        <xs:element name="users">
          <xs:complexType>
            <xs:sequence>
              <xs:element name="user" type="identity" minOccurs="0"
maxOccurs="unbounded"/>
            </xs:sequence>
          </xs:complexType>
        </xs:element>
      </xs:sequence>
    </xs:complexType>
  </xs:element>
  <xs:element name="groups">
```



```

<xs:complexType>
  <xs:sequence>
    <xs:element name="group" minOccurs="0" maxOccurs="unbounded">
      <xs:complexType>
        <xs:complexContent>
          <xs:extension base="identity"/>
        </xs:complexContent>
      </xs:complexType>
    </xs:element>
  </xs:sequence>
</xs:complexType>
</xs:element>
<xs:element name="members">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="member" minOccurs="0" maxOccurs="unbounded">
        <xs:complexType>
          <xs:sequence>
            <xs:element name="groupName" type="xs:string">
              <xs:annotation>
                <xs:documentation>Group LDAP DN</xs:documentation>
              </xs:annotation>
            </xs:element>
            <xs:element name="memberName" type="xs:string">
              <xs:annotation>
                <xs:documentation>Member LDAP DN</xs:documentation>
              </xs:annotation>
            </xs:element>
            <xs:element name="memberType">
              <xs:annotation>
                <xs:documentation>Group member is a "user" or "group"</
xs:documentation>
              </xs:annotation>
            </xs:element>
          </xs:sequence>
        </xs:complexType>
      </xs:element>
    </xs:sequence>
  </xs:complexType>
</xs:element>
<xs:complexType name="identity">
  <xs:sequence>
    <xs:element name="dn" type="xs:string">
      <xs:annotation>

```

```

    <xs:documentation>LDAP Distinguished Name</xs:documentation>
  </xs:annotation>
</xs:element>
<xs:element name="id" type="xs:string">
  <xs:annotation>
    <xs:documentation>Unique ID, typically LDAP object uid or cn attribute.
Defined per customer policies.</xs:documentation>
  </xs:annotation>
</xs:element>
<xs:element name="cn" type="xs:string">
  <xs:annotation>
    <xs:documentation>LDAP Common Name (CN) Attribute Value</xs:documentation>
  </xs:annotation>
</xs:element>
<xs:element name="description" type="xs:string" minOccurs="0">
  <xs:annotation>
    <xs:documentation>LDAP Description Attribute Value</xs:documentation>
  </xs:annotation>
</xs:element>
</xs:sequence>
</xs:complexType>
</xs:schema>

```

User Import Preview Transaction File

You can configure the Windows User Import, LDAP User Import, or XML User Import tools to generate a preview file, to see a sequential listing of the transaction it would perform if enabled for processing. Refer to “Using Windows User Import” on page 273, “Using LDAP User Import” on page 277, or “Using XML User Import” on page 281 for the procedure to generate this file. It is highly recommended that you generate and review a Preview Transaction File before importing users and groups. This is especially important if you are using this tool for the first time. In addition, review the dirsync.log for any potential errors. If numerous errors occur, they must be resolved prior to running the User Import tool.

Example Preview Transaction File

A brief example of a preview transaction file might look like this:

```
AGRP, Domain Users, All Domain Users, Admin, , , , , LKLINE3, , DEFAULT, , English,
YES,
UGRP, Domain Users, , , , , , LKLINE2, , , , ,
DGRP, Domain Users
AUSR, jsmith, John Smith, Manager, SRVDEVTEST\jsmith, Users, , YES, NO, NO, Admin, ,
,
AUGP, Domain Users, jsmith
DUGP, Domain Users, jjones
UUSR, jsmith, John Smith, Manager, , , , , , , , ,
DUSR, Ckline
```

Preview Transaction File Format

Each transaction entry begins with an operation code that identifies an action User Import would perform:

| | |
|------|--------------------------|
| AGRP | Add Group |
| UGRP | Update Group |
| DGRP | Delete Group |
| AUSR | Add Userid |
| UUSR | Update Userid |
| DUSR | Delete Userid |
| AUGP | Add Userid to Group |
| DUGP | Delete Userid from Group |

Field Definitions

Each operation has a fixed series of comma-delimited fields containing arguments. These are defined as follows:

- **Add Group (AGRP) operations**

```
AGRP, Domain Users, All Domain Users, Admin, , , , , LKLINE3, , DEFAULT, ,
English, YES,
```

| Field | Source |
|------------|------------------|
| Group Name | Domain directory |

| Field | Source |
|---------------------------------------|-------------------|
| Group Comment | Domain directory |
| Group Owner Name | Registry defaults |
| Group Default File Category | Registry defaults |
| Group Default Dispose Category | Registry defaults |
| Group Default Checkout File Name | Registry defaults |
| Group Default Checkout File Extension | Registry defaults |
| Group Primary Server | Registry defaults |
| Group Apps | Registry defaults |
| Group CSI Default Index | Registry defaults |
| Group Access Domain | Registry defaults |
| Group Language | Registry defaults |
| Group Secure Search | Registry defaults |
| Group CSI Threshold | Registry defaults |

- **Update Group (UGRP) operations**

UGRP, Domain Users, , , , , , LKLINE2, , , , , ,

Update Group operations have the same fields as Add Group operations, but only changed fields are filled in. A minus sign (-) in a field indicates an entry that will be removed.

- **Delete Group (DGRP) operations**

DGRP, Domain Users

Update Group operations have only one field: Group Name

- **Add Userid (AUSR) operations**

AUSR, jsmith, John Smith, Manager, SRVDEVTEST\jsmith, Users, , YES, NO, NO, Admin, , ,

| Field | Source |
|----------------------------|--------------------------------|
| Userid | Domain directory |
| Full User Name | Domain directory |
| Userid Comment | Domain directory |
| Userid simplified login ID | Computed from domain directory |

| Field | Source |
|---------------------------|-------------------------------|
| Default Group | Current active domain groupid |
| User Default Path | Registry defaults |
| User Privilege Add Items | Registry defaults |
| User Privilege Add Users | Registry defaults |
| User Privilege Add Groups | Registry defaults |
| User Owner Name | Registry defaults |
| User Status | Registry defaults |
| User Default Run App | Registry defaults |

- **Update Userid (UUSR) operations**

UUSR, jsmith, John Smith, Manager, , , , , , , , , ,

UUSR operations have the same fields as AUSR operations, but only changed fields are filled in. A minus sign (-) in a field indicates an entry that will be removed.

- **Delete Userid (DUSR) operations**

DUSR, jsmith

DUSR operations have only one field: Userid

- **Add Userid to Group (AUGP) operations**

AUGP, Domain Users, jsmith

AUGP operations have two fields:

Group Name
 Userid

These actions always occur after the userid has been added, so the userid is always valid in the library system.

- **Delete Userid from Group (DUGP) operations**

DUGP, Domain Users, jsmith

DUGP operations have two fields:

Group Name
 Userid

Managing Multiple Library Systems

Some sites may require more than one library system (that is, multiple Property Managers and associated Storage Managers) to manage their information. To determine the number of library systems your site will need as well as the overall configuration of servers and clients, consider the following:

- How many users will be registered?
- How many users will be active on the library system at any time?
- Are users dispersed over multiple geographically located sites?
- What types of applications are being used to generate the stored files?
- What versions of operating systems are being used?
- Do some work groups require customized configuration settings?
- Will all users require the ability to search for files based on file contents as well as descriptive properties?
- Can additional application servers be made available and what are the configurations (for example, disk sizes, CPU speed, available RAM)?
- Will users require the ability to replicate identical document, user, group, and property data into multiple systems?

Some sites may want to set up multiple library systems for customization and organizational reasons, rather than solely performance. Sites with departments that require different sets of custom properties or special configuration may prefer multiple library systems.

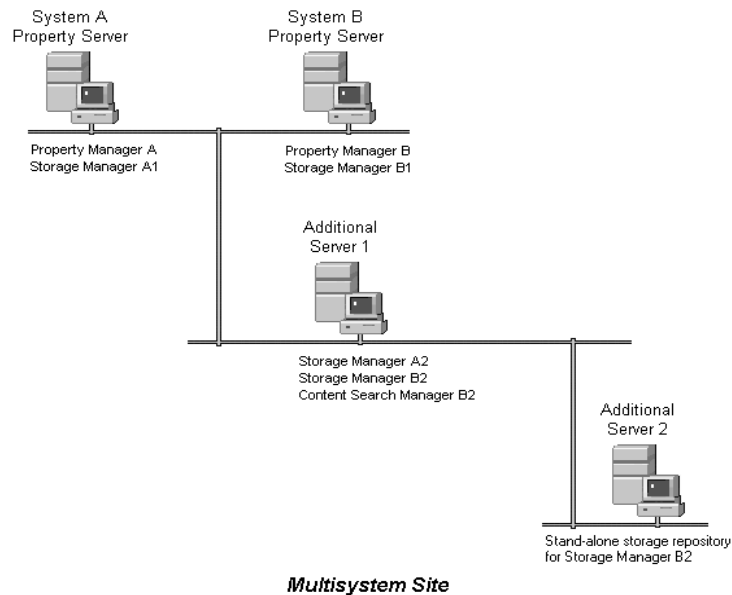
For example, if you have one large accounting department that requires its own set of fifteen custom properties, you will be left with only five of these properties for all other customization needs. In this case, you might consider setting up an entirely separate library system for the accounting department and adding the remainder of the users to another library system.

Note Refer to the “[Installation Guide](#)” on page 76 for information on how to install a library system and to “[Moving Content Services](#)” on page 352 for details on how to move a library system to a different server.

Sharing Servers and Disk Space

Storage Managers can share the same server with other Storage Managers, as long as each Storage Manager is associated with a different library system.

The following illustration shows some of the common configuration features of a multisystem site:



In this illustration, systems A and B each have an additional Storage Manager, both of which have been installed on the same remote server. Also on that remote server, a Content Search Manager has been installed for system B.

Storage repositories assigned to a Storage Manager can be redirected to other machines, thus spreading the disk space requirements across the network.

In the illustration above, one of the storage repositories created for Storage Manager B2 has been placed on a network server. Thus, assigning remote storage repositories is one way of expanding the disk space available for storing library system controlled versions, especially if you cannot dedicate more machines to library system services.

If you install an optional Content Search Manager, you must take into account its disk space requirements, especially in multisystem configurations.

Optimizing Search Queries

The CS library system provides better performance for particular kinds of searches. This is due partly to the requirements of content search functionality, and partly to optimization of the search processing that contains Keyword property search criteria.

Using Indexed Properties

The object properties in the following lists are indexed in the library system's database tables. To improve the response time for searches, you should build the search criteria with indexed properties whenever possible. Search performance will improve regardless of which search operators are used. For most of these objects, the indexed properties include, but are not limited to, those properties that make up the object's key (uniquely identifying an object or object instance).

Those objects having more than one indexed property use a composite index for the listed properties:

| Object | Property Name |
|----------------------------------|--|
| Access Control Object | Parent Object Name User or Group Name User or Group Type Parent Object Type |
| Application Object | Application Name Parent Object Parent Name Application Name |
| Archive Category Object | Archive Category |
| Archive Repository Object | Archive Repository ID |
| Checkout Object | Item ID Version ID Checked Out By User |
| Content Search Repository Object | Content Search Repository ID |
| Content Search Manager Object | Server |
| Custom Object | ISV ID Object Type Object Key |
| Custom Property Object | Property Number Property Name |

| Object | Property Name |
|------------------------------------|---|
| Default Item Access Control Object | Source Object Type Source Object Name User or Group Type User or Group Name Object Type |
| Group Object | Group Name |
| Group Member Object | Group Name User Name |
| Item Object | Item ID Date Added Current Version Date Added By User Added By Group Original File Name Current Version Author All custom properties |
| Keyword Object | Item ID Keyword |
| Message Object | Message ID |
| Other System Object | System Name |
| Property Manager Object | System Name |
| Search Object | Search Name Search Type Search Reference |
| Session Object | Session Number |
| Storage Category Object | Storage Category |
| Storage Repository Object | Storage Repository ID |
| System Object | System Name |
| User Icon Object | User Name Object Type Object Key |
| User Object | User Name Network logon Name |
| Version Object | Version ID Item ID |

Custom properties can be indexed, depending upon the setting for the *Indexed?* property in the instance of the Custom Property object.

The LIKE Operator

Avoid using the LIKE operator and, wherever possible, use the following search operators instead:

- EQUALS, as in:

```
Added By User = John Doe
```

- GREATER THAN / LESS THAN combinations, as in:

```
Date Added > January 1, 1997
```

and

```
Date Added < March 1, 1997
```

Use wildcard characters with the LIKE operator only as follows:

- Use a wildcard character only at the end of a value on which you are searching (for example, John%). This LIKE search will be nearly as fast as an EQUAL search.
- Avoid beginning a value with a wildcard character (for example, %son).
- Avoid beginning and ending a value with a wildcard character (for example, %son%).

Using Secure Searches

If the Secure Searches? property is set to Yes in the System object, a user or group must have at least Viewer access rights to all the results returned for any search. This requirement means that the library system must do access checking as part of the search procedure, which slows response time.

Therefore, unless your site specifically needs secure searches (or unless its users often run content searches), we suggest you disable this feature by setting the Secure Searches? property value to No in the System object. If needed, you can then use the Secure Searches property of the User, Group, or Search object to enable secure searches in more restrictive, or highly specific cases.

With secure searches *enabled*, content search performance may actually improve, especially on system-wide searches for a particular word or phrase. This is because the number of files searched would be limited to those to which the user has at least Viewer access rights.

Specifying Keywords in Searches

- Avoid creating a complex set of search criteria statements connected by ANDs and ORs.
- Do not overuse particular Keyword property values (avoid assigning and searching on keywords that describes a large percentage of your items).

Using Content Search Criteria

In any search containing *content* search criteria, try to make the associated *property* search criteria as specific as possible to limit the number of results returned.

Remember, however, that when contents and property search criteria are combined, the property search tends to be the faster portion of the search.

When only content values (and no property values) are specified for a search, an optimized search algorithm is applied. Depending on the environment, a content-only search may be a viable choice for best performance.

Content Search Filters

When a request is submitted to index a file, the Content Search Manager determines the file's format by examining the contents of the file. If a filter is available for the format, the file is converted to ASCII. If a file is recognized as being in text format already, it will be indexed as is (that is, as if it were already ASCII). If the format is not recognizable or a filter is not available, the file is not indexed.

Notes

- If you need to index a file type not on this list, and if you have a filter for that file type, contact FileNet Professional Services for assistance in setting it up for use.
- The KeyView filter for indexing Microsoft Word files does not index any words in headers or footers of files when the filename extension is DOC or RTF. In case you want to index headers or footers in such files, You need to change your *style.uni* as following:

```
type: "application/msword"
  /format-filter = "flt_kv -headfoot"
  /charset      = guess
  /def-charset  = 1252
```

```
type: "application/rtf"
  /format-filter = "flt_kv -headfoot"
  /charset      = guess
  /def-charset  = 1252
```

Archive Formats

| Format | Versions |
|--------|---------------|
| PKZIP | through 2.04g |

Computer-Aided Design

| Format | Versions |
|--------------------------|----------------------|
| AutoCAD Drawing | R13, R14, 2000, 2004 |
| AutoCAD Drawing Exchange | R13, R14, 2000, 2004 |
| Microsoft Project | 98, 2000, 2002 |

| Format | Versions |
|-----------------|------------------------|
| Microsoft Visio | 5, 6, 2000, 2002, 2003 |

Display Formats

| Format | Versions |
|--------------------------------|------------|
| Adobe Portable Document Format | 1.1 to 1.6 |

Graphic Formats

| Format | Versions |
|----------------------------------|----------------|
| Computer Graphics Metafile | N/A |
| CorelDRAW (TIFF header) | through to 9.0 |
| DCX Fax System | TIFF/CCITT/DCX |
| Encapsulated PostScript (raster) | TIFF header |
| Enhanced Metafile | N/A |
| Graphic Interchange Format | 87, 89 |
| JPEG | N/A |
| Lotus AMIDraw Graphics | N/A |
| Lotus Pic | N/A |
| Macintosh Raster | 2 |
| MacPaint | N/A |
| Microsoft Office Drawing | N/A |
| PC PaintBrush | 3 |
| Portable Network Graphics | N/A |
| SGI RGB Image | N/A |
| Sun Raster Image | N/A |
| Tagged Image File | 3.0 to 6.0 |
| Truevision Targa | 2 |
| Windows Animated Cursor | N/A |
| Windows Bitmap | N/A |
| Windows Icon Cursor | N/A |

| Format | Versions |
|-----------------------|----------|
| Windows Metafile | 3 |
| WordPerfect Graphics1 | 1 |
| WordPerfect Graphics2 | 2, 7 |

Mail Formats

| Format | Versions |
|--|----------------------|
| Microsoft Outlook | 97, 2000, 2002, 2003 |
| Microsoft Outlook Express | N/A |
| Microsoft Outlook Personal Folder (Windows only) | 97, 2000, 2002, 2003 |

Multimedia Formats

| Format | Versions |
|-----------------------|----------------------|
| MPEG-1 Audio Player 3 | ID3 versions 1 and 2 |

Presentation Formats

| Format | Versions |
|------------------------------|----------------------|
| Applix Presents() | 4.0, 4.2, 4.3, 4.4 |
| Corel Presentations | 7, 9, 10, 11, 2000 |
| Lotus Freelance Graphics | 96, 97, 98, R9, 9.8 |
| Lotus Freelance Graphics2 | 2 |
| Microsoft PowerPoint Windows | 97, 2000, 2002, 2003 |
| Microsoft PowerPoint Windows | 95 |
| Microsoft PowerPoint PC | 4 |
| Microsoft PowerPoint Mac | 98 |

Spreadsheets Formats

| Format | Versions |
|---------------------------|---------------|
| Applix Spreadsheets | 4.2, 4.3, 4.4 |
| Comma Separated Values () | N/A |

| Format | Versions |
|-----------------------------|-------------------|
| Coral Quattro Pro | 5, 6, 7, 8 |
| Lotus 1-2-3 | 96, 97, R9, 9.8 |
| Lotus 1-2-3 | 2, 3, 4, 5 |
| Lotus 1-2-3 Charts | 2, 3, 4, 5 |
| Microsoft Excel Macintosh | 98 |
| Microsoft Excel Windows | 2.2, through 2003 |
| Microsoft Works for Windows | 1.x to 4.0 |
| Microsoft Excel Charts | 2, 3, 4, 5, 6, 7 |
| Microsoft Works Spreadsheet | 1, 2, 3, 4 |

Text and Markup Formats

| Format | Versions |
|--------------------------------------|---------------|
| ANSI | N/A |
| ASCII | N/A |
| HTML | 3, 4.0 |
| MIME | N/A |
| IBM DCA/RFT (Revisable Form Text) | SC23-0758-1 |
| Microsoft Excel Windows XML | 2003 |
| Microsoft Word Windows XML | 2003 |
| Microsoft Visio XML | 2003 |
| OpenOffice | 1, 1.1 |
| Rich Text Format | 1 through 1.7 |
| StarOffice | 6, 7 |
| Unicode Text | 3, 4 |
| XHTML | 1.0 |
| XML (generic) | 1.0 |

Word Processors Formats

| Format | Versions |
|---|-----------------------------|
| Adobe Maker Interchange Format | 5, 5.5, 6, 7 |
| Applix Words | 3.11, 4, 4.1, 4.2, 4.3, 4.4 |
| DisplayWrite | 4 |
| Folio Flat File | 3.1 |
| Fujitsu Oasys | 7 |
| JustSystems Ichitaro | 8, 9, 10, 12 |
| Lotus AMI Pro | 2, 3 |
| Lotus AMI Professional Write Plus | 2.1 |
| Lotus Word Pro (Windows x86 only) | 96, 97, R9 |
| Lotus SmartMaster (Windows x86 only) | 96, 97 |
| Microsoft Word PC | 4, 5, 5.5, 6 |
| Microsoft Word Windows | 1.0 and 2.0 |
| Microsoft Word Windows | 6, 7, 8, 95 |
| Microsoft Word Windows | 97, 2000, 2002, 2003 |
| Microsoft Word Macintosh | 4, 5, 6, 98 |
| Microsoft Works | 1, 2, 3, 4 |
| Microsoft Works | 6, 2000 |
| Microsoft Windows Write | 1, 2, 3 |
| WordPad | through 2003 |
| WordPerfect Windows | 5, 5.1 |
| WordPerfect Windows | 6, 7, 8, 9, 10, 11, 2000 |
| WordPerfect Linux | 6, 8 |
| WordPerfect Macintosh | 1.02, 2, 2.1, 2.2, 3, 3.1 |
| XyWrite | 4.12 |

Configuring Content Search Stopwords

Stopwords can be configured in following ways:

Collection specific

Using style.stp file created in <install drive>\filenet\<library>\content\index\<repository>\style one can create a collection specific stopword list. This file lists the words that should not be indexed. Words on this list do not make it into a collection's word index, and therefore are not searchable. The style.stp supports the following:

- Regular expressions
- Case-sensitive matching (not case-insensitive)
- Only single-byte character sets

Locale specific

A stopword file for each language (locale) is located as shown in the following table:

| OS Platform | Location |
|-------------|--|
| Windows | InstallDrive:\filenet\Verity\common\language\vdk30.stp |
| UNIX | \$IDMDS_HOME/Verity/common/language/vdk30.stp |

Modify the vdk30.stp file to add or delete locale-specific words to/from the list.

Note FileNet recommends that stopwords not to be used for content searches. If you need a stopword file, consult the Verity K2 documentation for information.

Function-based Indexes

If your database engine is the Enterprise Edition of Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0, function-based indexes are available within Oracle when performing case-insensitive searches. Function-based indexes improve search performance only for case-insensitive searches, and should not be used unless performance without such indexes is poor. Moreover, function-based indexes require queries to be based on cost-based optimization, not rule-based optimization. You can set up function-based indexes only by performing a procedure after installing (or upgrading to) Content Services 5.5. You cannot set up these indexes during installation of Content Services or dynamically at runtime.

To enable the cost-based optimizer to use function-based indexes, set the following Oracle session or system variables:

```
QUERY_REWRITE_ENABLED=TRUE
QUERY_REWRITE_INTEGRITY=TRUSTED
COMPATIBLE=8.1.0 or higher
```

The command `QUERY_REWRITE_INTEGRITY=TRUSTED` needs to be in effect only while building the function-based index. You can place the `QUERY_REWRITE_ENABLED=TRUE` command in the file **init.ora**.

After creating the function-based index you need to run the ANALYZE command on the index and the table containing the column on which the index is built. The statistics you specify in the ANALYZE command can be gathered with either the ESTIMATE or the COMPUTE option.

For example, use the following SQL statement to create a function-based index named `index1` on a column named `col1` in table named `table1` (where the function `NLS_UPPER` returns a string all of whose letters are uppercase):

```
CREATE INDEX index1 ON table1 (NLS_UPPER(col1));
```

The following procedure provides a complete example of how to create function-based indexes for a library `lib1`. The procedure refers to some important function-based indexes that can be created. Some sites may need indexes on additional columns not referenced in the procedure.

To create a function-based index:

1. Choose the function-based indexes you will need.

2. Edit the `init.ora` file as follows:

- Set `QUERY_REWRITE_ENABLED` to `TRUE`
- Set `COMPATIBLE` to `8.1.0` or higher

3. Restart the Oracle instance.

4. Start a `SQL*PLUS` session and run the following session command:

```
ALTER SESSION SET QUERY_REWRITE_INTEGRITY=TRUSTED;
```

5. Create all the function-based indexes. The commands listed below are examples illustrating how to create function-based indexes and are recommended for Content Services. The next-to-last command in the list, creating the function-based index `fe_ina03`, is for an indexed single-valued custom property. The last command in the list, creating the function-based index `fitmv01`, is for an indexed multivalued custom property. All the others are for ordinary (non-custom) properties. You may want to create additional function-based indexes for *string* custom properties that are indexed.

```
CREATE INDEX fi3840_2 ON access_control(NLS_UPPER(ac_inst_nm));
CREATE INDEX fi256_4 ON element(NLS_UPPER(e_org_filename));
CREATE INDEX fi256_5 ON element(NLS_UPPER(e_icon_title));
CREATE INDEX fi512_1 ON version(NLS_UPPER(v_file_name));
CREATE INDEX fi8960_4 ON custom(NLS_UPPER(co_key));
CREATE INDEX fe_ina03 ON element(NLS_UPPER(e_ina03));
CREATE INDEX fitmv01_value ON item_mv_custom01(NLS_UPPER(itmv01_value));
```

6. Collect statistics on each library table and for all indexes, including the function-based indexes. The two commands listed below are examples illustrating how to collect statistics for the table `access_control` and a function-based index `mustangl.1256_4`:

```
ANALYZE TABLE lib1.access_control COMPUTE STATISTICS;
ANALYZE INDEX mustangl.1256_4 COMPUTE STATISTICS;
```

The following procedure uses a plan table to verify that the optimizer is using a function-based index that you have created for library `LIB1`.

To verify that a function-based index is being used:

1. Create the plan table by running the following command:

```
oracle\rdbms\admin\utlxplan.sql
```

2. Run the following command to generate an explain plan for a SQL statement that selects e_name and e_org_filename from the lib1.element table where the e_org_name is (case-insensitive) 'file.txt':

```
EXPLAIN PLAN SET STATEMENT_ID = 'verify' FOR
SELECT e_name, e_org_filename FROM lib1.element
       WHERE NLS_UPPER(e_org_filename) = NLS_UPPER('file.txt');
```

3. To display the explain plan, run the following command:

```
SELECT LPAD(' ', 2*(level-1)) || operation, options,
       object_name, position)
FROM plan_table
START WITH id = 0 AND statement_id = 'verify'
CONNECT BY prior id = parent_id AND statement_id = 'verify';
```

The operation column of the explain plan output should show that an index is being used in the query (for example, an index range scan).

Audit Logging

The Audit Logging feature allows administrators to capture data on all usage-related library system events. Library system auditing can capture all processes that users run from their desktop, regardless of the application. This means that auditing can capture IDM Desktop, Web Services, and CS Explorer user activities.

The audit data may help to determine potential attempts to breach library system security. Using the audit log information, you can:

- Audit suspect activities.
- Monitor specific library systems and the activities within them.
- View the audit trail by accessing the database tables directly, using your own SQL report generating software.

The audit log includes information on:

- Who initiated an action.
- The client workstation used.
- The date and time of the action.
- The specific action taken.
- The object acted upon (if any).
- The result of the action (success or failure).
- Additional information specific to the operation performed.

Notes

Audit Logging monitors only library system events. Events such as document printing are not under the control of library system services software, and therefore cannot be monitored by library system Audit Logging.

The library system auditing feature is different from the auditing features provided by Oracle and MS-SQL Server. The databases' auditing features capture SQL events; library system auditing captures document management events and stores them in Audit Event objects, which can be logged to provide an immediately useful record of how users are interacting with library system.

The library system database schema documentation is available in the [“Creating Reports Using Database Schemas” on page 462](#) section of this document.

For tool-specific instructions on using Audit Logging, refer to the *CS Explorer Online Help*.

Audit Logging Features

With the library system Audit Logging feature you can:

- Turn auditing on and off globally.
- Turn individual events on and off.

- Log events that succeed.
- Log events that fail.
- Never log an event.
- Clear some or all of the audit log based on date or set number.

The library system Audit Logging feature monitors these events:

| | |
|------------------------|------------------|
| Add Object | Suspend Session |
| Get Object | Resume Session |
| Delete Object | Add Item |
| Modify Object | Checkin Version |
| Check Object Existence | Checkout Version |
| Get Object Access Info | Cancel Checkout |
| Rename Object | De-index Version |
| Modify Password | Archive Version |
| Logon | Index Version |
| Logoff | Reclaim Version |
| Change Group | Search |

You can specify which audit events to monitor by modifying the Audit Events object. For tool-specific instructions on using Audit Logging, refer to the *CS Explorer Online Help*.

Each time a user performs an audited operation, an audit log entry is made. Because the audit log can grow quite large, it is best to carefully select the events to audit for your library system.

Note When any of the audit settings are changed, the changes affect subsequent user sessions. Any currently active user sessions (including the administrative session) do not reflect the audit changes.

Setting Up and Managing Auditing

When Audit Logging is enabled, the database can grow quite large. You may need to allocate additional space for the database (and for the transaction log). The amount of disk space required depends on:

- Which user events you choose to audit.
- The total number of user events audited.
- How many users are actively using the system.
- How long you plan to keep the audit log online before clearing it.

After turning auditing on, you should monitor the growth of the system for a few days to get an idea of your additional space requirements.

Auditing is constrained by database row and column size limits. The event data that is logged is truncated if the data captured for that event exceeds the set limit. The audit feature appends an ellipsis at the end of the strings that were truncated.

Turning Auditing Logging On or Off

Audit Logging can be turned on or off for an entire library system, and individual audit event items can be set to log on success, to log on failure, or to not log at all. For instructions on how to turn auditing on and off, refer to the *CS Explorer Online Help*.

Viewing the Audit Log

You can view the audit log from CS Explorer, or you can use SQL utilities to create your own view of the audit log.

- For instructions on viewing the audit log, refer to the *CS Explorer Online Help*.
- For an explanation of the audit log entries, refer to “[Audit Log File Format](#)” on page 325.
- For instructions on clearing audit log entries and creating audit log sets, refer to “[Managing the Audit Log File](#)” on page 327.
- For instructions on creating audit log reports using SQL utilities, refer to “[Audit Log Reports](#)” on page 329.

Audit Log File Format

The Audit log data is formatted as follows.

| Log Item | Description |
|--------------------------|--|
| Audit Event ID | The specific event which occurred. |
| Object ID | (Optional) The type of object acted upon. |
| Object Key | (Optional) The unique identifier that differentiates the object acted upon. |
| Param 1 | Operation specific parameters and values. Refer to “ Event Parameters ” on page 326 for details. |
| Param 2 | A list of operation specific parameters and values. Refer to “ Event Parameters ” on page 326 for details. |
| Param 3 | A list of operation specific parameters and values. Refer to “ Event Parameters ” on page 326 for details. |
| Param 4 | A list of operation specific parameters and values. Refer to “ Event Parameters ” on page 326 for details. |
| Search Request Start Key | (Searches only) The start key. |
| Search Request Direction | (Searches only) Search direction. |
| Search Request Count | (Searches only) The hit count. |

| Log Item | Description |
|--------------|---|
| Date Time | The date and time of the action. |
| User | The user who initiated the action. |
| Status | The status code, where 0 = success. |
| Workstation | The network node name of the client workstation where the action originated. |
| Elapsed Time | Duration of the operation, in milliseconds. |
| Set Number | The logical set number which this set belongs to. For more information on audit log sets, refer to “ Managing the Audit Log File ” on page 327. |
| Process ID | The ID of the client process which initiated the action. |
| Flags | (Optional) Flags specified for the operation. |
| Filename | (Optional) The source or destination filename passed to the operation. |

Event Parameters

Library system functions that add or modify objects create event parameters in the audit log. The event parameters are stored as a parallel list of attributes and values. Parameters 1 and 3 contain a list of each attribute name, separated by commas. Parameters 2 and 4 contain the corresponding values, separated by commas.

The parallel list is structured like this example:

| | | | | |
|---------|-------------------------------|---------------|---------------------------------|-----|
| Param 1 | Current Version Checked Out?, | Protected?, | Allow Secondary Version Lines?, | ... |
| Param 2 | No, | Yes, | Yes, | ... |
| Param 3 | Permanent Version?, | Checked Out?, | Check In By User, | ... |
| Param 4 | No, | No, | georgew, | ... |

This information is displayed in the audit log as follows:

| | |
|---------|--|
| Param 1 | Current Version Checked Out?,Protected?,Allow Secondary Version Lines?,... |
| Param 2 | No,Yes,Yes,... |
| Param 3 | Permanent Version?,Checked Out?,Check In By User,... |
| Param 4 | No,No,georgew,... |

To minimize the size of the Audit Log, event parameters **may not** contain duplicate information if that information appears elsewhere. For example, information contained in the Object Key is not duplicated in the event parameter.

Sample Audit Log Entry

Audit Event - Add Item

| | |
|----------------|----------|
| Audit Event ID | Add Item |
|----------------|----------|

| | |
|--------------------------|---|
| Object ID | Item |
| Object Key | 003670007 |
| Param 1 | Current Version Checked Out?,Protected?,Allow Secondary Version Lines?,Storage Category,Archive Category,Default Checkout Filename,Online Limit,Online Version Control,Online Disk Space,Current Version ID,Current Version Author,Last Accessed... |
| Param 2 | No,Yes,Yes,Documents,Trash,7016.txt,5,1,1035,1,georgew,General Users,Yes,0,1 |
| Param 3 | Permanent Version?,Checked Out?,Check In By User,Checked In By Group,Last Access User,Last Access Group,Availability Status,Secondary Version Line Count,Merge-Source Version ID,Merge-Destination Version ID,File Size,Storage Repository,La... |
| Param 4 | No,No,georgew,General Users,georgew,General Users,Online,0,0,0,1035,1,Yes,Index Pending,No |
| Search Request Start Key | |
| Search Request Direction | |
| Search Request Count | |
| Date Time | 1999-07-04T09:55:57,607 |
| User | georgew |
| Status | 0 |
| Workstation | martha |
| Elapsed Time | 640 |
| Set Number | 44 |
| Process ID | 5911 |
| Flags | 16 |
| Filename | 7016.txt |

Managing the Audit Log File

The rate of audit log growth depends on library system usage and the number of events being audited. To effectively manage auditing, you can separate events into distinct sets, and delete old sets after carefully reviewing their contents. You establish the conditions that determine when a new set begins. Each audit log entry is stamped with the current set number when it is added.

For example, you may decide to group all events in a single day into a set and to increment the set number at 6:00 pm each working day. You can then review the set for the day while events are logged to the next set. After you have analyzed the events in a given set, you can delete the set to conserve disk space.

Alternatively, you can retain a set indefinitely and delete old entries based on their date. For example, you can delete log entries at the end of the month.

Note FileNet recommends that you enable Audit Logging for only those events which require monitoring and that you delete old audit log entries when no longer needed. Otherwise, the audit log may become unwieldy.

Content Services provides three stored procedures to delete audit log entries and to increment the set number:

- **stpTruncAuditLog**
- **stpTruncAuditByDate**
- **stpIncrementSet**

There is no CS Explorer interface for these stored procedures. You must run them from a standard SQL tool such as Query Analyzer or SQL Plus.

To copy the audit log data to a file, use a SQL utility such as **bcp**. For more information, refer to “[Audit Log Reports](#)” on page 329.

stpTruncAuditLog

This stored procedure deletes all audit entries having a set number less than or equal to the specified set number. The number of rows deleted will be displayed.

Any entries having a set number greater than the *current* set number will also be deleted. This situation would occur if the set number rolled over. See **stpIncrementSet** for more information.

| | |
|--------|---|
| MSSQL | If no set number is specified, the current set number is incremented and all previous audit entries will be deleted. |
| Oracle | A set number must always be specified. Use a set number of zero (0) to indicate that all current entries should be deleted and the set number incremented. No row count or messages will be returned. |

stpTruncAuditByDate

This stored procedure deletes all audit log entries prior to the specified date, and displays the number of rows deleted. To prevent simultaneously occurring transactions from being deleted, the specified date must be prior to the current date.

| | |
|--------|---|
| MSSQL | MS SQL Server allows a variety of date formats to be used for the date parameter. If you enter an incorrect date format, an error message is displayed. |
| Oracle | The actual parameter passed to Oracle is a string which is converted inside the stored procedure to a date in the format MM-DD-YY. |

stpIncrementSet

This stored procedure increments the current set number without taking any other action. Use this stored procedure if you wish to differentiate between groups of transactions without deleting any entries. At a later date, you can use **stpTruncadtTopAll** to delete the transactions in set order. To delete only entries with a given set number while retaining entries with lower set numbers, use direct SQL commands.

Audit Log Reports

You can view the audit log with CS Explorer, or you can use your database utilities to create reports. For smaller library systems running MSSQL Server, you can use the **bcp** utility to export the audit information into a spreadsheet for simple analysis. For a high-volume library system, you can use SQL database audit report utilities to analyze the log data.

To create the SQL view of the audit log:

Internally each audit log record contains only a numerical ID indicating the event represented. Content Services displays a text string referenced by this numerical ID. You can use the following database view to generate reports that display the more useful event description instead of a cryptic ID.

Define the database view using the following SQL statement. Enter all the text on one line: it is presented here on several lines for easier reading. The precise form of the SQL statement depends on your database.

For a SQL Server 2000 database, the SQL statement is as follows:

```
create view AuditLog as select
EVENT_ID = NLS_TEXT,
OBJECT_KEY = AL_OBJECT_KEY,
EVENT_PARAM1 = AL_EVENT_PARAM1,
EVENT_PARAM2 = AL_EVENT_PARAM2,
EVENT_PARAM3 = AL_EVENT_PARAM3,
EVENT_PARAM4 = AL_EVENT_PARAM4,
SR_START_KEY = AL_SR_START_KEY,
SR_DIRECTION = AL_SR_DIRECTION,
SR_COUNT = AL_SR_COUNT,
DATETIME = AL_DATETIME,
USR = AL_USER,
STATUS = AL_STATUS,
WORKSTATN_ADDR = AL_WORKSTATN_ADDR,
ELAPSED_TIME = AL_ELAPSED_TIME,
SET_NUMBER = AL_SET_NUMBER,
PROCESSID = AL_PROCESSID,
FLAGS = AL_FLAGS,
FILENAME = AL_FILENAME from nls_data,
audit_log where NLS_MODULE_ID = 63999 and
NLS_STRING_ID = AL_EVENT_ID + 64 and
NLS_LANGUAGE_ID = (select L_LANGUAGE from library)
```

For an Oracle database, the SQL statement is as follows (be sure to replace <LIBRARYNAME> with the library name):

```
create view <LIBRARYNAME>.AuditLog as select
NLS_TEXT "EVENT_ID",
AL_OBJECT_KEY "OBJECT_KEY",
AL_EVENT_PARAM1 "EVENT_PARAM1",
AL_EVENT_PARAM2 "EVENT_PARAM2",
AL_EVENT_PARAM3 "EVENT_PARAM3",
AL_EVENT_PARAM4 "EVENT_PARAM4",
AL_SR_START_KEY "SR_START_KEY",
AL_SR_DIRECTION "SR_DIRECTION",
AL_SR_COUNT "SR_COUNT",
```

```
AL_DATETIME "DATETIME",
AL_USER "USR",
AL_STATUS "STATUS",
AL_WORKSTATN_ADDR "WORKSTATN_ADDR",
AL_ELAPSED_TIME "ELAPSED_TIME",
AL_SET_NUMBER "SET_NUMBER",
AL_PROCESSID "PROCESSID",
AL_FLAGS "FLAGS",
AL_FILENAME "FILENAME" from <LIBRARYNAME>.nls_data,
<LIBRARYNAME>.audit_log where
NLS_MODULE_ID = 63999 and
NLS_STRING_ID = AL_EVENT_ID + 64 and
NLS_LANGUAGE_ID = (select L_LANGUAGE from
<LIBRARYNAME>.library);
```

You can customize the display order of the column names, but the referenced column names and the rest of the view definition syntax must be exactly as shown in the SQL statements above. Do not insert carriage returns in the view definition.

To use the bcp utility:

Change to the *Drive:\Mssql\Binn* directory and enter the following command:

```
bcp librarysystem..AuditLog out filename -U login_id -P password -N
```

where:

librarysystem is the library system name

filename is the output filename for the report

login_id is the system administrator login name

password is the system administrator password for the database

Secure Document Delete

The Secure Document Delete feature, when enabled, will scrub a version file on the storage server during a delete operation, as well as during such common operations as automatic or ordinary archiving, reclaim daemon operations, and indexing. When documents are removed using Secure Document Delete, they are deleted permanently from those repository directories with no possibility of recovery via disk recovery tools.

Scrubbing means that the file's entire content is overwritten by a byte pattern before the file is actually deleted.

Secure Document Delete does not delete files from client machines. To ensure secure delete functionality on servers containing library system components, do not install client interfaces (such as IDM Desktop) on those servers.

Neither Replication Services nor Web Services scrub files during add, checkin, and checkout operations. For simple delete operations as may be desired during a “purge” process however, Secure Document Delete is applied on all servers that have this feature turned on.

Using Secure Document Delete

You can configure the Secure Document Delete feature by selecting one of the options described below.

Refer to “[Storage Manager Tab](#)” on page 167 for configuration procedures.

| Level | Description |
|-------------------|---|
| No Secure Deletes | Ordinary delete. Files are not overwritten before being deleted. |
| One Scrub | Files are overwritten with zeroes once before being deleted. This level corresponds to the Clear definition of the DOD 5220.22-M specification. |
| Three Scrubs | Files are overwritten once with an arbitrary character, again by the character's complement, and finally by a random character before being deleted. This level corresponds to the Purge or Sanitize definition of the DOD 5220.22-M specification. |

The initial and additional (remote) Storage Managers of a library system are each separately configurable with respect to secure deletes. Whether an item (document) is securely deleted depends on the secure-delete setting of all the Storage Managers that govern the storage repositories where the versions of the item are stored.

When an item is added to a library, it is given a Storage Category, which is associated with the storage repository where the versions of the item will be stored. As one storage repository fills up, the system administrator will assign the Storage Category to another storage repository. Thus, over time, different versions of the same item may end up in different storage repositories.

To ensure that an item is securely deleted, make sure it is given a Storage Category all of whose associated storage repositories will only be on servers where the Storage Manager is configured for secure deletes. Beware that if even one of the storage repositories associated with the Storage Category is on a server whose Storage Manager is not configured for secure deletes, then any versions of the item that are stored on that server will not be securely deleted.

Which Files are Scrubbed

Secure deletes are carried out only in the places that move or delete files being managed by a Storage Manager, or subsets of information about these files. For example:

- When delete item is called, the target file is scrubbed.
- When versions are archived from the storage repositories, the source is scrubbed.
- When lost versions are moved into the lost+found directory, the source is scrubbed.
- When files are indexed any temporary copy files, and filtered files (temporary or otherwise) that have completed their lifetime are scrubbed.
- The reclaim daemon scrubs the **rcmd**, **rfile**, and **dfile** directories during a reclaim of a version file from offline storage.

Which Files are Not Scrubbed

Library files created on, copied to, or left behind on a client machine are not controlled by Secure Document Delete. Files will be insecurely deleted on the client, for example, during checkin. You must manage the secure deletion of these files using your own tools.

Turning on Secure Document Delete does not clean up directories containing files that have already been deleted.

When files are moved for archiving, the source file is scrubbed. Once files are archived, however, their secure deletion from archive repositories (**dfile** directories to which archived document versions are moved) is the responsibility of the system administrator. That is, you must use your own tools to securely delete those files as they permanently archive, reclaim, or dispose of them from archive repositories.

During de-index operations, the entry is marked for deletion from the index, but the indexed words are not removed until the next merge occurs.

Content Search Repositories going through indexing rebuilds or deletes are not scrubbed. Administrators must delete securely the old index using their own tools prior to rebuilds.

During a content search indexing merge (a regular internal operation), the indexing software does not scrub its own temporary files, including the old immediate file during the shrinking of this file. However, such files are encoded in binary form and the contents are not discernible.

Any file contained on a backup is insecure and is not controlled by Secure Document Delete. You must manage the long-term secure retention of these backups using your own tools.

The REPCHECK utility, after checking repository consistency with the library system database, securely moves orphaned files to the lost+found directory. You must use your own tools to securely delete those files as you permanently archive, re-add, or dispose of them. Because of possible loss of consistency between the files in the repositories and the records in the database, we strongly recommend that you run the REPCHECK utility periodically. Refer to [“REPCHECK” on page 439](#) for more information.

Running the csindex, csmerge, or csfilter executables directly does not scrub the temporary files unless they are run with the correct command-line parameters. Contact FileNet Technical Support for assistance in doing this.

Secure Document Delete Limitations

On UNIX platforms, files larger than 2 Gbytes will not be overwritten and must be scrubbed using special tools. On Windows platforms, files will be scrubbed securely up to the full Windows 64-bit limit.

Secure Document Delete does not scrub property data in the database; this remains the job of the database administrator and the database itself. Microsoft SQL Server and Oracle do not scrub database data during row or table deletion.

You cannot obtain Secure Document Delete functionality on compressed disks or files. Do not use defragmentation tools on the disk containing the stored document files. Such tools move the files on the disk without pre-scrubbing them, making secure deletes unsupported.

Some database metadata (such as file names, Item IDs, Version IDs) are not scrubbed in subtle places. For example, when a search query is performed, even on the server, a temp file is used to store the results. When the query ends, the temp file is deleted (but not securely). Database metadata is generally insecure; even Microsoft SQL Server and Oracle databases cannot do secure metadata deletes.

The Windows pagefile and UNIX swap files may include memory images that need to be paged to disk, and the memory image may include the files you are copying or writing. The Secure Document Delete functionality cannot scrub these files.

Reformatting a hard disk may not overwrite data on the disk. If you reformat a disk containing storage repositories, you should use other tools to scrub the disk clean of file data.

The uninstall command (**dsuninst.exe**) does not scrub storage or index repository files.

Activities such as the actions of the executable files that perform during installs and upgrades, the use of tools, and utilizing configuration code, do not carry out secure deletes. For example when DLLs are copied to a temp directory, they are not scrubbed when they are removed.

Secure Document Delete does not work on Hierarchical Storage Manager extended drives. HSM controls access to files on media and cannot guarantee scrubbing in the process of a move or purge.

Deleting Sessions

In a library system, an active session is the open session that a user is currently working in. There are actually two sessions involved when a user is logged in to a library system via a Content Services client application, such as IDM Desktop or CS Web Admin environment. One is the database connection session, and the second is the Content Services library session.

When a CS client application exits or is killed, the corresponding database connection session is deleted from the database. You can also delete all the database connections at once by stopping and restarting the database. If a database connection session is deleted, then any corresponding active CS library session is said to be an orphan session.

CS library system sessions are tracked by CS Explorer but cannot be removed by it. Neither stopping and restarting the database or CS services will delete an orphan session. Nor will rebooting a workstation where a CS client application is running.

To remove these orphan CS library sessions, you can invoke DSSTOP on the server where the library system is resident. Alternatively, you can simply allow the daemon running on the database custodial server to periodically check for suspended sessions. The daemon will remove any orphaned sessions that are more than two days old.

To find out what sessions are open in a particular library system and to see whether these sessions are active or suspended, you can check the status of the Sessions object using CS Explorer.

Note For tool-specific instructions on how to view the status of a session, refer to the *CS Explorer Online Help*.

Resetting the Session ID

The library system maintains a unique session ID for each library system session. The session ID is incremented each time a user opens a new session. Over time, it is possible to exceed the limit of approximately 2 billion (2^{31}) session IDs.

To avoid an overflow condition, we recommend that you periodically check the value of the Session Number property in a Session object (the value of this property is the session ID). When the value of the Session Number property approaches 2 billion, you should reset the session ID to 0.

To reset the session ID:

1. Stop the library system and database services.
2. Restart the database services.
3. Using your SQL query tool (**SQL Server Query Analyzer** for SQL Server or **sqlplus** for Oracle), reset the session ID to 0.

- For SQL Server enter:

```
use system_name
go
delete from session
go
update numid set se_id_num=0
go
```

- For Oracle enter:

```
delete from system_name.session;  
update system_name.numid set se_id_num=0
```

4. Restart the library system services.

Backing Up the Library System

As a library system administrator, you are responsible for backing up the library systems using your own backup and restore tools. From the backups, you can restore a library system in the event that:

- Hardware failure occurs on one or more library system server.
- The property repository becomes corrupted.
- Files in any of the storage repositories become lost or inaccessible.
- A user loses a version file.
- The content search repositories are corrupted or lost.

To accomplish a safe level of backup, you will need to create and regularly schedule automated backup routines that back up the following information:

- The database components that contain library system property information.
- Version files and the names and locations of version files stored in the storage repositories.
- All storage repository directories.
- Content search repository files and index data.
- All reclaimable archive repository directories.

Note To maintain correct synchronization of replicas and their source objects under all circumstances (including failure of a library system server), be sure to schedule backup operations so that all library systems participating in replication are backed up at the same time. This ensures that you will be able to synchronize restore operations on all these library systems, and thereby maintain cross-system data integrity. Refer to [“Replication Services” on page 593](#) for more information.

Planning the Backup

For each library system, identify and make note of all of its components and where they reside. There are 5 main questions you need to answer before you create backup scripts:

1. For each library system, how many Storage Managers, Archive Managers, and Content Search Managers are there?
2. Where are these components? What are the names of the servers on which they reside?
3. Which component is located on which server?
4. What are the specific file locations and directory structures of those components. (Use CS Explorer to find the various directories.)
5. What is your backup strategy (daily, full, incremental, differential). Remember that your backup strategy should be based on your restore process.

Using this information, you can now write your backup script files, which you should schedule to be run regularly.

Housekeeping

You may want to perform the following database housekeeping functions at the same time as your scheduled database backup:

- Use the REPCHECK utility to report and fix inconsistencies between information in the property repository and the storage and content search repositories. Refer to “REPCHECK” on page 439 for more information.
- Use your own database tools to run a consistency check against the property repository. (MS-SQL users can use the *dbcc* facility.)
- Create a log of the files being backed up and store the log in the library system as an indexed file (to enable a content search for it). Then, when you need to restore a file:
 - a. Determine the encrypted file name using CS Explorer. Each Version object has an Encrypted File Name property that indicates the name under which the file is stored on the storage repository.
 - b. Search the library system for the file name. The search should return the log file that identifies the backup saveset where the file is located.

Backing Up Property Repositories

Use the tools supplied with your database engine (or whatever third-party tools you have available) to perform regular backups of your library system's property repository. Many library system administrators do daily incremental backups and weekly full backups.

- Microsoft SQL Server installations must include both the database and its transaction logs whenever doing a full backup. An incremental backup is simply a backup of the transaction log.
- Oracle installations wishing to perform incremental, or partial, backups must operate in ARCHIVELOG mode. Because an archived redo log is present, the data files restored from a partial backup can be made consistent with the rest of the database during recovery procedures.

Backing Up Storage Repositories

Use your own tape backup software to back up the storage repository files located in the following directories.

| Platform | Location |
|----------|---|
| Windows | <i>Drive:\filenet\LibrarySystem\Stacks\Shelfxxx</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/stacks/shelfxxx</i> |

Backing Up Archive, Reclaim, and Reclaim Command Repositories

Use your own tape backup software to back up the archive repository files located in the following directories.

| Platform | Location |
|----------|---|
| Windows | <i>Drive:\filenet\LibrarySystem\Stacks\dfilexxx</i> |
| | <i>Drive:\filenet\LibrarySystem\Stacks\rfilexxx</i> |
| | <i>Drive:\filenet\LibrarySystem\Stacks\rcmdxxx</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/stacks/dfilexxx</i> |
| | <i>\$IDMDS_HOME/librarysystem/stacks/rfilexxx</i> |
| | <i>\$IDMDS_HOME/librarysystem/stacks/rcmdxxx</i> |

Backing Up Content Search Repositories

Use your own tape backup software to back up the files located in the following directories. (These directories exist only if you have installed a Content Search Manager.)

| Platform | Location |
|----------|---|
| Windows | <i>Drive:\filenet\LibrarySystem\Content\Index</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/content/index</i> |

Make sure the **csmerge** utility is not scheduled to run at the same time as your backup. The Hour of Day to Cleanup property in the Content Search Repository object controls when **csmerge** is run.

Note During the time specified by the Hour of Day to Cleanup property, and whenever **csmerge** is running, no content-based retrieval or indexing takes place.

Note A scheduled **csmerge** execution will run only if **csmerge** has not been run within the preceding 24 hours.

Restoring Content Services

At some point, you may need to restore one or more components of a library system due to corrupted data or a device or server failure.

Here are some typical situations that would call for you to restore library system components:

- The database engine or the property repository (database) are installed on a device that fails.
- The database engine and the property repository are installed on a server that fails.
- A storage repository or a content search repository are on a device or server that fails.
- A library system service (Storage Manager, Content Search Manager, or Replication Services) is on a device or a server that fails or a content search repository is installed on a server that fails.
- A repository or a library system service becomes corrupted.

Restoring library system components—data files (repositories) or software (library system services)— may involve retrieving data files from backup and/or reinstalling software. The restore procedures apply to situations where you are restoring data files and reinstalling software on the server where they were originally installed (on the same device or a different device) or on another server with the same name. If the new location for the data files and software is on a server whose name differs from that of the original server, refer to the procedures in [“Moving Content Services” on page 352](#).

Caution For a library system that participates in replication, FileNet strongly recommends that you always restore a library system service to a server with the same name as the server where the component was originally installed. If you need to restore the library system service to a server with a different name, contact [FileNet Worldwide Customer Support](#) for instructions.

If you need to restore the database, do so before restoring any other repositories or library system services. To fully recover from the loss of one or more components of a set of library services (Storage Manager, Content Search Manager, or Replication Services) caused by hardware or software failure, you must reinstall the library services software (the CS executable files) and then locate and manually restore the repository files that were lost from the associated repositories. To recover from the loss of a property, storage, archive, or content search repository caused by hardware or software failure, you will have to locate and manually restore the repository files that were lost.

- To restore the database, refer to [“Restoring the Database” on page 339](#)
- To restore library services, refer to [“Restoring Library Services” on page 340](#).
- To restore storage repositories, refer to [“Restoring Storage Repositories” on page 343](#).
- To restore archive repositories, refer to [“Restoring Archive Repositories” on page 345](#).
- To restore content search repositories, refer to [“Restoring Content Search Repositories” on page 346](#).

Restoring the Database

The following procedure shows how to restore a database and/or database engine, either on its original server or on a different server with the same name.

To restore the database:

1. If the library system you want to restore is still operable:
 - a. Put the library system in single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
 - b. Shut down all library system services.
 - c. Log out of the library system.
2. If you are reinstalling the database engine, either on its original server or on a new server, refer to “Preparing an Oracle Database”.
3. Use your Oracle database administration tools to restore your property repository from backup.
4. If you are also restoring library services, continue at “Restoring Library Services” on page 340; otherwise, continue at step 5.
5. Start the library system services. Refer to “Controlling Library System Services” on page 198 for information about stopping, starting, and checking the status of services.
6. When you are ready to allow users to begin using the library system again, use CS Explorer to set Logons Enabled? to Yes.

Restoring Library Services

This section deals with how to restore a set of library services to its original server or to a new server with the same name. If you want to restore a set of library services to a new server with a different name, refer to “Moving Content Services” on page 352.

If you need to restore the database, do so before restoring library services (refer to “Restoring the Database” on page 339).

Caution Because losing the storage repositories on a library system server may potentially affect your most recently created and modified documents, we urge you to perform frequent system backups and to keep complete backup records on other media. Such precautionary actions can save you and your users hours of recovery time, and help you avoid the loss of your most valuable current data. For details on other steps you can take to prepare for a storage repository failure, refer to “Preparing for a Possible Repository Failure” on page 227.

Before performing the following procedure, choose the devices where you are going to restore the library services, replacing any corrupt devices as needed.

To restore a set of library services:

1. Run the following command to uninstall the Content Search Manager.

WARNING Before running `dsuninst -v`, back up the indexes for each set of library services you plan to uninstall. The `\content\index` directory will be removed during the Verity uninstall process.

```
DSUNINST -v library_name
```

2. If the library system whose services you want to restore is still operable, use CS Explorer to put the library system into single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
3. Log out of the library system.
4. On the database server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services are running.
5. On the server where you are reinstalling CS executable files, shut down any applications that use ODBC to connect to the database, such as Oracle Enterprise Manager.
6. On the server where you are reinstalling CS executable files, stop the services for all previously installed library systems.

Note If you prefer, you can perform [step 7](#) through [step 20](#) noninteractively (refer to “[Reinstalling CS Executable Files Unattended](#)” on page 364).

7. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the server where you are reinstalling CS executable files and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.
8. On the Welcome Screen, click Next.
9. On the Software License Agreement screen, click Approve, and then click Next.
10. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. Your database software should be Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0. Click Next.
11. On the Select Setup Operation screen, click Reinstall Executables and then click Next. Continue at [step 13](#) (in a non-cluster environment) or [step 12](#) (in a cluster environment).
12. On the Specify Virtual Server Name screen, enter the virtual server name, assigned by the domain administrator, that identifies the cluster where you are reinstalling executable files. Then click Next.
13. On the Specify Reinstall Information screen, enter the requested information and click Next:
 - a. If you have not changed the ODBC data source name on this server, do not modify the value shown. Otherwise, type the new value.
 - b. Type the name of the tablespace corresponding to the library services you are reinstalling.
 - c. Clear the check box.
14. On the Specify Verity Settings screen, specify the Verity K2 Master Admin server name and port (as needed) and the Verity locale (language), as described below, and then click Next.

- a. If the screen has no check box, skip to step b, since the library services you are reinstalling will be governed by an existing Master Admin (whose server name and port are shown on the screen) either on this server or some other server. If the screen does have a check box, do one of the following:
 - Select the check box and specify the server name and port of an existing Master Admin on some other server to govern the library services you are reinstalling.
 - Clear the check box to install a Master Admin (the server name and port are shown on the screen) with the library services you are reinstalling on the server.

Note If no previously installed set of library services remains on this server, and no library services are installed on any other server in your organization, you *must* clear the check box so that a Master Admin will be installed with the library services you are reinstalling on the server.

- b. The locale you choose (from the drop-down list) applies only to the Verity collections (indexes) on this server for this library system, and can differ from the locale for services of this library system on other servers. It can also differ from the locale for services of other library systems on this server.

15. In the Specify Install Locations screen, supply the requested information and click Next:

- a. CS Shared Files: the directory where the CS shared files are to be reinstalled is located. The suggested location is ***InstallDrive:\Program Files\FileNet\Shared***. (If this directory still exists from an earlier install of library services on this server, you cannot change it.)

For each of the following CS components, choose a partition with at least 20 MB of free space where the component will be reinstalled. A drop-down list for each component shows the free space in the available partitions. In a cluster environment, you need to choose a partition on a drive shared by all nodes in the cluster.

- b. General Services: the partition where miscellaneous CS utilities will be reinstalled.
- c. Storage Manager: the partition where the Storage Manager will be reinstalled.
- d. Content Search Manager: the partition where the Content Search Manager will be reinstalled.

Note For optimum performance, install the Content Search Manager in its own partition.

Note For General Services, choose a partition with at least 2.2 GB of free space.

After you click Next, CS Setup will prompt you to create directories for the CS shared files and other CS components if they don't exist. Click Yes at each prompt.

16. In the Specify Final Setup Values screen, supply the requested information and click Next.

- a. Installer Name: the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
- b. Installation Comment: the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is "Content Services 5.5".
- c. Select or clear the check box that determines if you want the services for this library system to automatically start whenever the server is rebooted.

- d. Select or clear the check box that determines if the services for this library system will automatically start when CS Setup exits. (This check box is available only if another set of library services is already installed on this server.)
17. The Review Setup Entries screen shows your reinstallation selections in the preceding steps. (You can also see your selections in the **setup.inf** file, located at *InstallDrive:\filenet*.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and reinstalling software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

18. In the Installing Components screen, review the results of the reinstallation and click Next. If the log on the screen indicates problems in the installation, refer to “[Problem Solving](#)” on page 215. If no set of library services, other than the set you are reinstalling, is on this server, skip to [step 20](#); otherwise continue at [step 19](#).
19. On the Setup Complete screen, clear the check box to omit performing another installation task within the current execution of CS Setup. Then click Next.
20. The appearance of the Exit Setup screen depends on whether you previously installed a set of library services on the server prior to the current execution of CS Setup:
 - a. If you previously installed services for another library system on the server, simply click Finish to exit CS Setup.
 - b. If you are reinstalling the only set of library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

21. Reinstall any Content Services patches that were installed prior to restoring the library services.
22. Start the library services using CS Manager (refer to “[Content Services Manager](#)” on page 198) if they are not already running. This will re-create the associated storage and archive repositories.
23. Continue at “[Restoring Storage Repositories](#)” on page 343.

Restoring Storage Repositories

If the property repository is not affected by the storage repository failure, you can rely on the property repository's records to help you replace the storage repository files for each restored storage repository, and to reconcile the differences between those restored files and the current data in the property repository.

If you know you have lost one or more storage repositories but do not know on which servers they are located, use CS Explorer to display the Storage Repositories property values for the System object, and check the Server property assignment for each storage repository.

Caution Because a lost storage repository may potentially affect your most recently created and modified items, we urge you to perform frequent system backups and to keep complete backup records on other media. Such precautionary actions can save you and your users hours of recovery time and help you avoid the loss of your most valuable current data. For details on other steps you can take to prepare for a storage repository failure, refer to “Preparing for a Possible Repository Failure” on page 227.

To restore a storage repository:

1. If the library system is still operable, put the library system in single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
2. If you need to restore a storage repository to a new device, and the old and new devices do not have the same identifier, use CS Explorer to edit the Device property value in the corresponding Storage Repositories object to correspond to the new device.
3. Copy the appropriate storage repository directories and files to the restored storage repository device. Using your own network tools, perform this step for each storage repository in one of the following ways:
 - a. If you are able to determine exactly which versions must be replaced, copy only these versions to the restored storage repository.
 - b. If you have no specific list of versions that must be replaced, copy all the most recently backed up versions associated with the restored storage repository. If the storage repository fills up during this copy operation, do the following:
 - i. Run the REPCHECK utility for the storage repository. Refer to “REPCHECK” on page 439 for instructions.
 - ii. Based on the report from REPCHECK, remove the versions that should not be in the storage repository.
 - iii. Copy as many more of the remaining versions to the storage repository as you can.
 - iv. Repeat steps i - iii until all the backed up versions have been copied into the storage repository.
4. We recommend that you reassign the Read-Only privileges for all restored storage repository files using the appropriate command for your operating system. The library system automatically marks the storage repository files as Read-Only when the files are created to prevent client applications from inadvertently changing or deleting these files.
5. Repeat [step 2](#) through [step 4](#) for any other storage repository you need to restore on this server.

6. If you are restoring library services in addition to restoring storage repositories, continue at “Restoring Archive Repositories” on page 345. If you are restoring storage repositories only, set the Logons Enabled? property value in the System object to Yes, and notify users that library system sessions can be opened again.

Restoring Archive Repositories

If you lose an archive repository, you will also lose the repository files that were awaiting final archiving after leaving the library system's control. This data loss may not be a problem if the archive repository is for nonreclaimable versions. However, if the lost versions were to be archived and possibly reclaimed by users, you must determine which archived versions were lost and how you can recover them.

Caution Because a lost archive repository may affect files that would potentially need to be reclaimed at a later date, we urge you to keep track of where archived versions are located in your backup storage medium, and archive as quickly as possible those versions that have been moved to reclaimable archive repositories.

The library system tracks the archived versions as far as the archive repository and records the archiving date (and time) in the library system. With this information, it may be possible for you to recover most of the lost files from your backup media, assuming that the archived versions (or similar earlier versions) were stored on that media. However, you will need a method of determining exactly which repository files must be replaced.

The following procedure enables you to restore a archive repositories and recover the repository files that were lost before they could be archived.

To restore archive repositories:

1. If the library system whose services you want to restore is still operable, use CS Explorer to put the library system into single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
2. If you are restoring an archive repository to a new device, and the old and new devices do not have the same identifier, use CS Explorer to edit the Device property value in the corresponding Archive Repositories object to correspond to the new device.
3. Determine which repository directories and files were lost with the archive repository, as follows:
 - a. Determine which documents may have had versions archived since you last copied versions from the archive repositories to the storage media.

Each Item object instance contains an Archive Category property that may help you make this preliminary determination, since archive categories are assigned to a particular archive repository.
 - b. For each document that may have had a version archived, check the Versions property of the Item object for the Archive Date value.
 - c. Determine which archived versions of each document were possibly on the archive repository when it was lost. The value of the Archive Repository ID property will identify the archive repository, and the value of the Archive Date property will provide the exact date and time when the version was placed in that archive repository.

4. Copy the archived versions that were lost with the archive repository to the archive or storage media as follows:
 - a. (Optional) Retrieve the versions that were copied to backup media before they were lost with the archive repository. Do this only if you want to store final archive copies of the versions in a different location from the backup copies. You can copy these files to the archive repository first or directly to your storage media.
 - b. Have users re-create the files for the archived versions that were not recoverable from other media. For example, users may have to re-create the archived versions that had never been backed up. Generally, users will need to modify other existing versions to recreate such files.
5. Repeat [step 2](#) through [step 4](#) for any other archive repository you need to restore.
6. If you are restoring library services in addition to restoring archive repositories, continue at “[Restoring Content Search Repositories](#)” on [page 346](#). If you are restoring archive repositories only, set the Logons Enabled? property value in the System object to Yes, and notify users that library system sessions can be opened again.

Restoring Content Search Repositories

Perform the procedure in this section to restore content search repositories. If the property repository is not affected by a content search repository failure, you can rely on the property repository's records to help you replace the storage repository files for each restored content search repository, and to reconcile the differences between those restored files and the current data in the property repository.

The steps in the procedure that deal with restoring Verity collections refer to a library system named *lib1* on a server named *server1*. Where these names appear in the procedure, substitute the name of your own library system and server.

To restore a content search repository:

1. If the library system whose services you want to restore is still operable, use CS Explorer to put the library system into single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
2. Using your own network tools, copy your **content\index** directory into the intended device location *install_drive:\filenet\lib1*.
3. If the intended device is different from the previous device, use CS Explorer to edit the Device property value in the corresponding Content Search Repositories object to correspond to the new device.
4. Run the CSK2CFG utility to register and attach the Verity components of the library system you are restoring.
5. Run the REPCHECK utility to verify that the logical links between the content search repository and the property repository. Respond to any discrepancies found in the REPCHECK utility report log. Refer to “[REPCHECK](#)” on [page 439](#) for instructions.
6. Stop and restart the library system services.

7. Set the Logons Enabled? property value in the System object to Yes and notify users that library system sessions can be opened again.

To determine whether the content search repository is correctly attached to its Verity K2 Server, and to attach it if it isn't, run the utility CSK2CFG. Refer to [“CSK2CFG” on page 386](#) for more information.

Renaming Servers

You can rename database servers, initial storage servers, and additional storage servers. However, the following restrictions apply:

- FileNet does **not** support renaming servers that are members of a cluster.
- For library systems that participate in replication, FileNet strongly recommends that you **do not** rename your database server.
- Renaming a database server will render all IDM Desktop stored searches and stored templates unusable.

Except where noted, each step in the following procedure applies to renaming database servers, initial storage servers, and additional storage servers.

To rename a server:

1. Change the host name of the server. (Refer to your Windows documentation for instructions.)
2. If you are using Active Directory to manage the domain where the server being renamed resides, use the ADSI Edit support tool on the primary domain controller to delete the old server name from the domain.
3. On the database server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services are running.
4. If you are renaming a database server,
 - a. Run the following SQL commands against the Master database:

```
sp_dropserver old_server_name;  
sp_addserver new_server_name, LOCAL;
```
 - b. On the database server, remove the net service name for the old server name and create a net service name for the new server by manually editing the **tnsnames.ora** file or by using Oracle Net Manager.
 - c. Perform step b also on every storage server where library services access the database server, and on every machine where a Content Services client application (such as IDM Desktop) accesses a library system supported by the database server.
 - d. On every machine where a Content Services client application (such as IDM Desktop) accesses a library system supported by the database server, use ODBC Data Source Administrator to remove the entry in the System Data Sources list for the old server name and add an entry for the new server name.

Perform the remaining steps of this procedure only if you are renaming a storage server. If you are renaming a database server that is *not* also a storage server, then skip the remaining steps.

5. If you are renaming a server where a Verity K2 Master Admin is installed, run the following SQL command once for each library system that has initial or additional services on the server:

```
update library set l_verity_admin='new_server_name';
```

The following steps reinstall the CS executable files for one set of library services. Perform these steps for each set of library services on the server being renamed. If a Verity K2 Master Admin is on the server being renamed, you

must also perform the steps for each set of library services—on all other storage servers—governed by this Verity K2 Master Admin.

6. Back up the storage, archive, and content search (Verity collections) repositories of the library system on the storage server (refer to “[Backing Up the Library System](#)” on page 336).
7. Using CS Admin Tools, check the Database Custodial Server property in the System object to see if server being renamed is the database custodial server. If it is, change the property value to the new server name.
8. On the server where you are reinstalling CS executable files, shut down any applications that use ODBC to connect to the database, such as Oracle Enterprise Manager.
9. On the server where you are reinstalling CS executable files, stop the services for all previously installed library systems.

Note If you prefer, you can perform [step 10](#) through [step 22](#) noninteractively (refer to “[Reinstalling CS Executable Files Unattended](#)” on page 364).

10. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the server where you are reinstalling CS executable files and double-click **Drive:\Setup.exe** to launch the Content Services Setup program. Alternatively, you can reinstall the CS executables noninteractively (refer to “[Reinstalling CS Executable Files Unattended](#)” on page 364) in place of performing [step 11](#) through [step 22](#).
11. On the Welcome Screen, click Next.
12. On the Software License Agreement screen, click Approve, and then click Next.
13. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. Your database software should be Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0. Click Next.
14. On the Select Setup Operation screen, click Reinstall Executables and then click Next.
15. On the Specify Reinstall Information screen, enter the requested information and click Next:
 - a. If you have not changed the ODBC data source name on this server, do not modify the value shown. Otherwise, type the new value.
 - b. Type the name of the tablespace corresponding to the library services you are reinstalling.
 - c. Clear the check box.
16. On the Specify Verity Settings screen, retain the existing Verity K2 Master Admin server name, port, and Verity locale (language), and then click Next.
17. In the Specify Install Locations screen, supply the requested information and click Next:
 - a. CS Shared Files: the directory where the CS shared files are to be reinstalled is located. The suggested location is **InstallDrive:\Program Files\FileNet\Shared**. (If this directory still exists from an earlier install of library services on this server, you cannot change it.)

For each of the following CS components, choose a partition with at least 20 MB of free space where the component will be reinstalled. A drop-down list for each component shows the free space in the available partitions. In a cluster environment, you need to choose a partition on a drive shared by all nodes in the cluster.

- b. General Services: the partition where miscellaneous CS utilities will be reinstalled.
- c. Storage Manager: the partition where the Storage Manager will be reinstalled.
- d. Content Search Manager: the partition where the Content Search Manager will be reinstalled.

Note For optimum performance, install the Content Search Manager in its own partition.

Note For General Services, choose a partition with at least 2.2 GB of free space.

After you click Next, CS Setup will prompt you to create directories for the CS shared files and other CS components if they don't exist. Click Yes at each prompt.

18. In the Specify Final Setup Values screen, supply the requested information and click Next.
 - a. Installer Name: the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
 - b. Installation Comment: the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is "Content Services 5.5".
 - c. Select or clear the check box that determines if you want the services for this library system to automatically start whenever the server is rebooted.
 - d. Select or clear the check box that determines if the services for this library system will automatically start when CS Setup exits. (This check box is available only if another set of library services is already installed on this server.)
19. The Review Setup Entries screen shows your reinstallation selections in the preceding steps. (You can also see your selections in the **setup.inf** file, located at *InstallDrive:\filenet*.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and reinstalling software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

20. In the Installing Components screen, review the results of the reinstallation and click Next. If the log on the screen indicates problems in the installation, refer to "Problem Solving" on page 215. If no set of library services, other than the set you are reinstalling, is on this server, skip to [step 22](#); otherwise continue at [step 21](#).

21. On the Setup Complete screen, clear the check box to omit performing another installation task within the current execution of CS Setup. Then click Next.
22. The appearance of the Exit Setup screen depends on whether you previously installed a set of library services on the server prior to the current execution of CS Setup:
 - a. If you previously installed services for another library system on the server, simply click Finish to exit CS Setup.
 - b. If you are reinstalling the only set of library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

After renaming a storage server, determine whether the Verity components on the renamed server are correctly configured, and repair the configuration if necessary, by running the utility CSK2CFG. Refer to “CSK2CFG” on [page 386](#) for more information.

Moving Content Services

At some point you may need to move the database (property repository) or library services from one server to another server, either because of a hardware/software failure or because you want to use a server with more resources. Or you may need to change the Windows domain membership of a storage server where a Verity K2 Master Admin is installed.

Caution

- If you are moving the database or are restoring a set of library services to the same server, or to another server with the same name, refer instead to the procedures in [“Restoring Content Services” on page 339](#).
- If you have a library system that participates in replication, and you need to move a database or a set of library services from one server to another server with a different name, contact [FileNet Worldwide Customer Support](#) for instructions.
- If you move your database to a server with a new name, all IDM Desktop stored searches and stored templates will be unusable.

To move the database, refer to [“Moving the Database” on page 352](#).

To move a set of library system services, refer to [“Moving Library System Services” on page 355](#).

To change the Windows domain membership of a storage server where a Verity K2 Master Admin is installed, refer to [“Moving a Verity K2 Master Admin” on page 360](#).

Moving the Database

Moving a database from one server to another involves more than just backing up the database on the old server and restoring it on the new server. You must also reinitialize the database user password (refer to [“Creating a Database User” on page 115](#)) for each library system that references the ODBC data source name for the old server. You accomplish this by reinstalling the Content Services executable files once for each library system that uses the database. You can run the reinstallation on any server where services for the affected library system are installed. To move the database from one server to another, perform the steps in the following procedure.

To move the database:

1. Back up the Oracle database on the old server.
2. Install on the new server the version of the Oracle database server software (Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0) that is on the old server.
3. Create the database and the database user (refer to [“Creating a Database User” on page 115](#)) on the new server. The name of the database on the new server should be the same as what it was on the old server.
4. Use CS Explorer to put all the library systems into single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.

- No other session but yours is open.
5. Log out of the library systems.
 6. Export the database schema for each library system on the old server and import them to the new server. Refer to “Exporting/Importing a Windows Oracle Database” on page 148 or “Exporting/Importing a UNIX Oracle Database” on page 150.
 7. On each storage server and on each machine running a CS client application (such as IDM Desktop, Web Services, or Open Client), change the ODBC data source to point to the new database server.

Note It is strongly recommended that you only modify the existing DSN, and only change the Server parameter, not the Name parameter, of the DSN. That way you won’t have to reconfigure your IDM Desktop-configured libraries; nor will you need to reinstall Content Services executable files.

You must now reinstall the Content Services executable files and reconfigure your IDM Desktop libraries, as shown in the remaining steps of this procedure. You need to reinstall only one set of library services for each library system that uses the database, no matter how many sets of library services are in the library system.

8. On the new database server, open the Windows Control Panel, double-click Administrative Tools, double-click the Services tool, and verify that the Oracle services are running.
9. On the server where you are reinstalling CS executable files, shut down any applications that use ODBC to connect to the database, such as Oracle Enterprise Manager.
10. On the server where you are reinstalling CS executable files, stop the services for all previously installed library systems.

Note If you prefer, you can perform [step 11](#) through [step 24](#) noninteractively (refer to “Reinstalling CS Executable Files Unattended” on page 364).

11. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the server where you are reinstalling CS executable files and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.
12. On the Welcome Screen, click Next.
13. On the Software License Agreement screen, click Approve, and then click Next.
14. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. Your database software should be Oracle9i 9.2.0 or Oracle 10g 10.2.0.1.0. Click Next.
15. On the Select Setup Operation screen, click Reinstall Executables and then click Next. Continue at [step 17](#) (in a non-cluster environment) or [step 17](#) (in a cluster environment).
16. On the Specify Virtual Server Name screen, enter the virtual server name, assigned by the domain administrator, that identifies the cluster where you are reinstalling executable files. Then click Next.
17. On the Specify Reinstall Information screen, enter the requested information and click Next:
 - a. If you have not changed the ODBC data source name on this server, do not modify the value shown. Otherwise, type the new value.

- b. Type the name of the tablespace corresponding to the library services you are reinstalling.
 - c. Select the check box and enter the temporary user password for the database user created in [step 3](#).
18. On the Specify Verity Settings screen, retain the existing Verity K2 Master Admin server name, port, and Verity locale (language), and then click Next.
19. In the Specify Install Locations screen, retain the existing installation locations and click Next.
20. In the Specify Final Setup Values screen, supply the requested information and click Next.
- a. **Installer Name:** the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
 - b. **Installation Comment:** the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is “Content Services 5.5”.
 - c. Select or clear the check box that determines if you want the services for this library system to automatically start whenever the server is rebooted.
21. The Review Setup Entries screen shows your reinstallation selections in the preceding steps. (You can also see your selections in the **setup.inf** file, located at *InstallDrive:\filenet*.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and reinstalling software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

22. In the Installing Components screen, review the results of the reinstallation and click Next. If the log on the screen indicates problems in the installation, refer to “[Problem Solving](#)” on [page 215](#). If no set of library services, other than the set you are reinstalling, is on this server, skip to [step 24](#); otherwise continue at [step 23](#).
23. On the Setup Complete screen, clear the check box to omit performing another installation task within the current execution of CS Setup. Then click Next.
24. The appearance of the Exit Setup screen depends on whether you previously installed a set of library services on the server prior to the current execution of CS Setup:
- a. If you previously installed services for another library system on the server, simply click Finish to exit CS Setup.
 - b. If you are reinstalling the only set of library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

25. For each library system using the database you have moved, run the DS_DSBL utility (*Drive:\filenet\library_system\pbin\ds_dsbl.exe*) to disable logons, as follows:

```
ds_dsbl database_server_name library_system_name
```

where the database server name is specified by the Oracle net service name.

26. Run the utility CAT_PROC (refer to “CAT_PROC” on page 377) to restore stored procedures.
27. On each server where you have IDM Desktop or Web Services installed, run the FileNet IDM Configure application (choose Start > Programs > FileNet IDM > Configure) to delete the configured libraries that reference the old Name parameter in the DSN and add the same libraries so that they reference the new Name parameter in the DSN. (The Name parameter corresponds to the “Host machine” parameter in the Configure application.)
28. Use CS Admin Tools to enable logons for each library system using the database that you moved.
29. Restart the library system services.

After moving the database to a new server, determine whether the Verity components on the new server are correctly configured, and repair the configuration if necessary, by running the utility CSK2CFG. Refer to “CSK2CFG” on page 386 for more information.

Moving Library System Services

To move a set of library services from one server, back up the CS data files on the old server, if the files are available. Then, run CS Setup on the new server, choosing the Reinstall Executables option, to install the library services on the new server. Finally, restore the CS data files on the new server and re-map all connections to the CS data files to point to their new locations.

Note It is highly recommended that you install the library system services in drive locations on the new server that are identical to those on the old server. For example, if the Storage Manager and Content Search Manager were respectively on the C: drive and D: drive of the old server, retain these drive assignments on the new server.

A Verity K2 Master Admin governs each set of library services (the Content Search Manager, in particular) for one or more library systems, including the library services that you are moving to a new server. As part of the installation of these library services on the old server, you either created a new Master Admin bundled with it (as part of Content Search Manager) or you “pointed” the library services at an existing Master Admin bundled with some other set of library services, either on the old server or on some other server.

To move library system services:

1. Put the library system in single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
2. Log out of the library system.

3. Back up all the storage, archive, and content search repositories on the old server.
4. (Recommended) Run an appropriate operating system command to verify access and data integrity of the locations where you intend to move library system software. Back up all the storage, archive, and content search repositories on the old server, if they are available.
5. On the database server, verify that the Oracle services are started.
6. On the server where you are reinstalling CS executable files, shut down any applications that use ODBC to connect to the database, such as Enterprise Manager or Query Analyzer.
7. The server to which you are moving a set of library system services may already have one or more previously installed sets of services for some other library systems. If so, execute SQL commands, as described in the following steps, to make sure that the set of library services you are moving has the same installation root as that of the library services already installed on the new server:

- a. Run the following SQL command to determine the installation root of a previously installed set of services for some library (say *lib1*) on the new server:

```
SELECT l_root FROM lib1.library;
```

- b. Run the following SQL command to determine the installation root of the set of services for the library (say *lib2*) you are moving to the new server:

```
SELECT l_root FROM lib2.library;
```

- c. If both SQL commands generate the same output then continue at [step 9](#), since the installation roots for both sets of library services are the same. If the installation roots are different, run the following SQL command to set the installation root of the *lib2* library services to the installation root for *lib1*. For example:

```
USE lib2
UPDATE dbo.library SET l_root='install_root1'
```

where *install_root1* is the installation root for the services of *lib1* on the new server.

8. If you are moving a set of library services (say, for library *lib2*) for which Replication has been enabled
 - a. Run the following SQL command:

```
USE lib2
UPDATE CUSTOM SET CO_STRING1='new_server_name', CO_STRING2='new_server_name'
where CO_ISV_ID='SAROSOPG' AND CO_OBJECT_TYPE='REPLCFG' AND CO_KEY='REPLCFG';
```

9. If you are moving a set of library services (say, for library *lib2*) that contains a Verity Master Admin,

- a. Run the following SQL command:

```
USE lib2
UPDATE library SET l_verity_admin='new_server_name'
```

- b. The following command uninstalls the Content Search Manager for one set of library services. On every storage server, run the command once for each of the other sets of library services governed by the Verity Master Admin being moved.

WARNING Before running **dsuninst -v**, back up the indexes for each set of library services you plan to uninstall. The `\content\index` directory will be removed during the Verity uninstall process.

```
DSUNINST -v library_name
```

- c. Restore each set of library services whose Content Search Manager you uninstalled in step b. Refer to “Restoring Library Services” on page 340.
10. On the server where you are reinstalling CS executable files, stop the services for all previously installed library systems.

Note If you prefer, you can perform [step 11](#) through [step 24](#) noninteractively (refer to “Reinstalling CS Executable Files Unattended” on page 364).

11. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the server where you are reinstalling CS executable files and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.
12. On the Welcome Screen, click Next.
13. On the Software License Agreement screen, click Approve, and then click Next.
14. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. Your database software should be SQL Server 2000. Click Next.
15. On the Select Setup Operation screen, click Reinstall Executables and then click Next. Continue at [step 17](#) (in a non-cluster environment) or [step 16](#) (in a cluster environment).
16. On the Specify Virtual Server Name screen, enter the virtual server name, assigned by the domain administrator, that identifies the cluster where you are reinstalling executable files. Then click Next.
17. On the Specify Reinstall Information screen, enter the requested information and click Next:
 - a. Type the new value for the ODBC data source name of the database.
 - b. Type the name of the tablespace corresponding to the library services you are reinstalling.
 - c. Clear the check box.
18. On the Specify Verity Settings screen, specify the Verity K2 Master Admin server name and port (as needed) and the Verity locale (language), as described below, and then click Next.
 - a. If the screen has no check box, skip to step b, since the library services you are reinstalling will be governed by an existing Master Admin (whose name and port are shown on the screen) either on this server or some other server. If the screen does have a check box, do one of the following:
 - Select the check box and specify the server name and port of an existing Master Admin on some other server to govern the library services you are reinstalling.
 - Clear the check box to install a Master Admin (the name and port are shown on the screen) with the library services you are reinstalling on the server.

Note If the library services you are moving to this server are the only services installed on any server in your organization, you *must* clear the check box so that a Master Admin will be installed with the library services you are reinstalling on the server.

- b. The locale you choose (from the drop-down list) applies only to the Verity collections (indexes) on this server for this library system, and can differ from the locale for services of this library system on other servers. It can also differ from the locale for services of other library systems on this server.

19. In the Specify Install Locations screen, supply the requested information and click Next:

- a. CS Shared Files: the directory where the CS shared files are to be reinstalled is located. The suggested location is ***InstallDrive:\Program Files\FileNet\Shared***. (If this directory still exists from an earlier install of library services on this server, you cannot change it.)

For each of the following CS components, choose a partition with at least 20 MB of free space where the component will be reinstalled. A drop-down list for each component shows the free space in the available partitions. In a cluster environment, you need to choose a partition on a drive shared by all nodes in the cluster.

- b. General Services: the partition where miscellaneous CS utilities will be reinstalled.
- c. Storage Manager: the partition where the initial Storage Manager will be reinstalled.
- d. Content Search Manager: the partition where the initial Content Search Manager will be reinstalled.

Note For optimum performance, install the Content Search Manager in its own partition.

Note For General Services, choose a partition with atleast 2.2 GB of free space.

After you click Next, CS Setup will prompt you to create directories for the CS shared files and other CS components if they don't exist. Click Yes at each prompt.

20. In the Specify Final Setup Values screen, supply the requested information and click Next.

- a. Installer Name: the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
- b. Installation Comment: the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is "Content Services 5.5".
- c. Select or clear the check box that determines if you want the services for this library system to automatically start whenever the server is rebooted.

21. The Review Setup Entries screen shows your reinstallation selections in the preceding steps. (You can also see your selections in the **setup.inf** file, located at ***InstallDrive:\filenet***.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and reinstalling software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

22. In the Installing Components screen, review the results of the reinstallation and click Next. If the log on the screen indicates problems in the installation, refer to “Problem Solving” on page 215. If no set of library services, other than the set you are reinstalling, is on this server, skip to [step 24](#); otherwise continue at [step 23](#).
23. On the Setup Complete screen, clear the check box to omit performing another installation task within the current execution of CS Setup. Then click Next.
24. The appearance of the Exit Setup screen depends on whether you previously installed a set of library services on the server prior to the current execution of CS Setup:
 - a. If you previously installed services for another library system on the server, simply click Finish to exit CS Setup.
 - b. If you are reinstalling the only set of library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

25. Restore the storage, archive, and content search repositories you backed up in [step 3](#) to the new server.
26. Start the CS services and then use CS Admin Tools to update the Server and Device properties in the storage, archive, and content search repository objects to reflect the new server name and device.
27. If the old server was the database custodial server for the library system whose services you have moved to a new server, use CS Admin Tools to change the Database Custodial Server property value to the name of the new server.
28. Run the CSK2CFG utility to register and attach the Verity components of the library system you are moving.
29. Reboot the new server and restart all library services on the new server.
30. Restart each Content Search Manager for the library system whose services you have moved.
31. Using CS Admin Tools, delete the Storage Manager and/or Content Search Manager on the old server.
32. On all machines running CS client applications that access the library system whose services you moved, reconfigure your connection to the library. For example, use the IDM Desktop Configure utility (available via the Windows Start menu) to delete and re-add the library system server name.
33. For every storage repository and content search repository that you have moved, run the REPCHECK utility to report on any inconsistencies between the database and the storage repositories or content search repositories. Refer to “REPCHECK” on page 439.
34. Check to see whether REPCHECK moved any files to the **lost+found** directory (located one directory up from storage repository directory). The library system administrator must take the appropriate action to address each file that appears in **lost+found**. For example, you may need to add the file to the appropriate storage repository (making the necessary modifications to the property repository) or you may need to delete the file.
35. For each library system from which users can open sessions to the library system whose services you are moving, use CS Admin Tools to delete any Other Systems object associated with the old server and add an Other Systems object associated with the new server name.

To determine whether the Verity components on the new server are correctly configured for the library services you moved, and to repair the configuration if necessary, run the utility CSK2CFG. Refer to “CSK2CFG” on page 386 for more information.

Moving a Verity K2 Master Admin

If you change the Windows domain membership of a storage server where a Verity K2 Master Admin is installed (the *Master Admin server*), you have, in effect, “moved” the Verity K2 Master Admin. And you must then also perform the following tasks for each library system governed by that Verity K2 Master Admin:

- Update the LIBRARY database schema table.
- Uninstall the Verity software for each set of library services governed by the Verity K2 Master Admin.
- Reinstall each set of library services governed by the Verity K2 Master Admin.

The following procedure shows the detailed steps involved in moving a Verity K2 Master Admin. In the procedure, assume that the Master Admin server is named *server1* and you wish to change its Windows domain membership from *old_domain* to *new_domain*.

To move a Verity K2 Master Admin:

1. Put each library system governed by the Verity K2 Master Admin in single-user mode. A library system is in single-user mode when the following conditions are met:
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.
2. Log out of each library system governed by the Verity K2 Master Admin.
3. Back up all the storage, archive, and content search repositories of each library system governed by the Verity K2 Master Admin.
4. Using Windows tools, change the Windows domain of the Master Admin server from *old_domain* to *new_domain*.
5. Run the following SQL commands for each library system whose services are governed by the Verity K2 Master Admin you moved to a new Windows domain in [step 4](#):

```
UPDATE libraryname.library SET l_verity_admin='server1.new_domain';
```
6. For each library system governed by the Verity K2 Master Admin you moved to a new Windows domain in [step 4](#), run the following command on every server where a set of initial or additional services for that library system is installed:

```
DSUNINST -v libraryname
```

For example, you may have a library system with its set of initial services on one server, and three sets of additional services on three other servers. In this case, you will run the DSUNINST command four times for this library system.

7. Run the following command once for each server that has library services governed by the Verity K2 Master Admin that you moved in [step 4](#):

```
dsuninst -v
```

You must now perform the following steps (which reinstall the CS executable files) once for each set of initial or additional library services governed by the Verity K2 Master you moved in [step 4](#).

8. On the database server, open the Windows Control Panel, double-click Administrative Tools, double-click Services, and verify that the Oracle services is started.
9. On the server where you are reinstalling the CS executable files, stop the services for all installed library systems.

Note If you prefer, you can perform [step 10](#) through [step 23](#) noninteractively (refer to “[Reinstalling CS Executable Files Unattended](#)” on page 364).

10. *FileNet Content Services* is packaged in four CD-ROMs. Insert the first *FileNet Content Services* CD-ROM into the CD-ROM drive of the server where you are reinstalling CS executable files and double-click **Drive:\Setup.exe** to launch the Content Services Setup program.
11. On the Welcome Screen, click Next.
12. On the Software License Agreement screen, click Approve, and then click Next.
13. The Before You Install screen reminds you to make sure that you have already installed database software that is certified for use with Content Services. Your database software should be SQL Server 2000. Click Next.
14. On the Select Setup Operation screen, click Reinstall Executables and then click Next. Continue at [step 16](#) (in a non-cluster environment) or [step 15](#) (in a cluster environment).
15. On the Specify Virtual Server Name screen, enter the virtual server name, assigned by the domain administrator, that identifies the cluster where you are reinstalling executable files. Then click Next.
16. On the Specify Reinstall Information screen, enter the requested information and click Next:
 - a. Type the value for the ODBC data source name of the database.
 - b. Type the name of the tablespace corresponding to the library services you are reinstalling.
 - c. Clear the check box.
17. On the Specify Verity Settings screen, specify the Verity K2 Master Admin server name and port (as needed) and the Verity locale (language), as described below, and then click Next.
 - a. If the screen has no check box, skip to step b, since the library services you are reinstalling will be governed by an existing Master Admin (whose name and port are shown on the screen) either on this server or some other server. If the screen does have a check box, select it and specify the server name and port of the Master Admin that you have moved to a new Windows domain.
 - b. The locale you choose (from the drop-down list) applies only to the Verity collections (indexes) on this server for this library system, and can differ from the locale for services of this library system on other servers. It can also differ from the locale for services of other library systems on this server.
18. In the Specify Install Locations screen, supply the requested information and click Next:

- a. CS Shared Files: the directory where the CS shared files are to be reinstalled is located. The suggested location is ***InstallDrive:\Program Files\FileNet\Shared***. (If this directory still exists from an earlier install of library services on this server, you cannot change it.)

For each of the following CS components, choose a partition with at least 20 MB of free space where the component will be reinstalled. A drop-down list for each component shows the free space in the available partitions. In a cluster environment, you need to choose a partition on a drive shared by all nodes in the cluster.

- b. General Services: the partition where miscellaneous CS utilities will be reinstalled.
- c. Storage Manager: the partition where the initial Storage Manager will be reinstalled.
- d. Content Search Manager: the partition where the initial Content Search Manager will be reinstalled.

Note For optimum performance, install the Content Search Manager in its own partition.

Note For General Services, choose a partition with atleast 2.2 GB of free space.

After you click Next, CS Setup will prompt you to create directories for the CS shared files and other CS components if they don't exist. Click Yes at each prompt.

19. In the Specify Final Setup Values screen, supply the requested information and click Next.
 - a. Installer Name: the name (up to 32 characters) that CS Explorer will display for the Installer property in the System object. The suggested name is Administrator.
 - b. Installation Comment: the text (up to 70 characters) that CS Explorer will display for the Comment property of the System object. The suggested text is "Content Services 5.4".
 - c. Select or clear the check box that determines if you want the services for this library system to automatically start whenever the server is rebooted.
 - d. Select or clear the check box that determines if the services for this library system will automatically start when CS Setup exits. (This check box is available only if another set of library services is already installed on this server.)
20. The Review Setup Entries screen shows your reinstallation selections in the preceding steps. (You can also see your selections in the **setup.inf** file, located at ***InstallDrive:\filenet***.) If you need to make changes, return (by clicking Back) to each screen controlling the selection you wish to change. Make your changes and return to the Review Setup Entries screen. Then click Next to begin copying files and reinstalling software.

After CS Setup program has finished copying files from the first *FileNet Content Services* CD-ROM, it would prompt you to insert the second CD-ROM.

- a. Insert the second CD-ROM.
- b. Browse to the location where disk07 folder is located and click OK.

Note Do not browse in disk07 folder.

Repeat the above steps for third and fourth CD-ROMs. For third CD-ROM, you need to browse to disk08 folder and for fourth CD-ROM disk09 folder.

21. In the Installing Components screen, review the results of the reinstallation and click Next. If the log on the screen indicates problems in the installation, refer to “[Problem Solving](#)” on page 215. If no set of library services, other than the set you are reinstalling, is on this server, skip to [step 23](#); otherwise continue at [step 22](#).
22. On the Setup Complete screen, clear the check box to omit performing another installation task within the current execution of CS Setup. Then click Next.
23. The appearance of the Exit Setup screen depends on whether you previously installed a set of library services on the server prior to the current execution of CS Setup:
 - a. If you previously installed services for another library system on the server, simply click Finish to exit CS Setup.
 - b. If you are reinstalling the only set of library services on the server, click the appropriate option button to either restart the computer now (automatically, as soon as CS Setup exits) or later (manually). If you choose to restart the computer now, remove the Content Services CD-ROM from the CD-ROM drive. In either case, click Finish to exit CS Setup.

Note Always remove the Content Services CD-ROM from the CD-ROM drive before restarting the computer.

24. Restart the Storage Manager, Content Search Manager, and Replication Services (if enabled) for the set of library services you have just reinstalled.
25. After reinstalling each set of services for a library system governed by the Verity K2 Master Admin that you moved to a new Windows domain, use CS Admin Tools to enable logons for the library system.

To determine whether the Verity components on the server are correctly configured in its new domain, and to repair the configuration if necessary, run the utility CSK2CFG. Refer to “[CSK2CFG](#)” on page 386 for more information.

Reinstalling CS Executable Files Unattended

As an alternative to running the interactive Content Services Setup program, you can run an unattended (silent non-interactive) Content Services installation program to reinstall initial or additional library system services.

To perform an unattended reinstallation, you need to do the following:

1. Prepare a parameter file to provide input to the Content Services installer program (refer to “[Preparing a Parameter File](#)” on page 364).
2. Copy the complete installation media to a local physical drive.
3. Invoke the Content Services installer program, using the parameter file as input to the program (refer to “[The dssetup Command Syntax and Options](#)” on page 138 for command options and syntax).

The Content Services installer program used for unattended installation of initial library services, additional library system services, and display languages is named **dssetup.exe** and is on the first *FileNet Content Services* CD-ROM.

Preparing a Parameter File

The parameter file providing input to **dssetup** can be generated in either of two ways.

- Make a copy of a sample parameter file from the first *FileNet Content Services* CD-ROM and edit it to meet the needs of your organization. Refer to “[Preparing a Parameter File](#)” on page 364.
- Run the (interactive) Content Services Setup program as far as the Review Setup Entries screen. But click Cancel, instead of Next, to exit from CS Setup without reinstalling any software. The parameter file containing your installation selections is **setup.inf**, located at *InstallDrive:\filenet*.

The sample parameter file **sample_reinstall.inf** for unattended reinstallation of initial or additional library services is available on the first *FileNet Content Services* CD-ROM.

The LibraryName entry in the parameter file (refer to “[Parameter File Format and Content](#)” on page 139) is case sensitive when you are reinstalling library services. Run the following SQL*Plus command to determine the exact case of your library name (for example, *mylib1*):

```
select l_name from mylib1.LIBRARY;
```

Then edit the LibraryName entry with the output of the SELECT command above. If the exact case of the library name is *myLib1*, then the LibraryName entry appears as follows:

```
LibraryName=myLib1
```

Using a Parameter File to Reinstall Library Services

Once you have prepared a parameter file, as explained above, you are ready to run **dssetup** to perform a reinstallation of initial or additional library services.

To perform an unattended library system reinstallation:

1. At a DOS command prompt invoke the **dssetup** command to perform the reinstallation. For example, if the sample file is named **setup.inf**, run the following command:

```
dssetup -f C:\setup.inf -g NO
```

This command will perform a reinstallation according to the parameters specified in **C:\setup.inf** without prompting you for inputs. Refer to “[Parameter File Format and Content](#)” on page 139 for command-line syntax and options.

2. Refer to “[Starting and Monitoring Library Services](#)” on page 128 for information about stopping, starting, and checking the status of library services.

Examples of Parameter Files

The following files are examples of parameter files that you can edit according to the needs of your organization to reinstall Content Services components. (Note that the `IndexTableName` entry in the file applies only when you are installing Content Services on an Oracle platform).

Example parameter file to reinstall initial or additional library services:

```
[setup]
Procedure=R
Component=S, T, R
VcatServerName=DBServer
VServerName=localhost
SystemName=libraryname
IndexTableName=
DBPassword=
RootName=filenet
InstDrive=C:
SharedDir=C:\Program Files\FileNet\Shared
SystemLanguage=
AdministrativeUser=
AdministrativePassword=
AdministrativeGroup=
UniversalGroup=
StorageCategory=
ArchiveCategory=
DocClass=
StorageRepositoryDrive=C:
IndexRepositoryDrive=C:
VerityMasterAdminServer=K2AdminServer.domain.local
VerityMasterAdminPort=9550
VerityIndexLanguage=English
Installer=Administrator
Comment=Content Services 5.5
AutoStart=0
StartNow=0
Verbose=1
LogFile=C:\filenet\cssetup.log
```

Document Classes and CVLs

Each document in a library system can be of a specific type or *Document Class*. For example, contracts or engineering specifications might each be a different Document Class in your library system. You can assign each Document Class its own set of standard and custom properties. A Document Class of Contracts might have a custom property for the lawyer's name, while a Document Class of Engineering might have a custom property for the project name. Other standard library properties, such as Title and Added by User, might be properties of several or all of your Document Classes. Once a Document Class is assigned to a document by a user, it can be changed by an administrator using the Content Services Admin Tools or by users running IDM Desktop.

Caution (FNDTS00071803) The total number of characters in the names of the properties applied to a document class must not exceed 256. If you exceed the limit, the extra characters are ignored, and you will not be able to log onto the library system to which the document class belongs. To recover from this error condition, you must remove as many properties from the document class as needed to drop below the 256-character limit.

Controlled Vocabulary Lists (CVLs) allow you to set up a list of valid entries for a custom property. You can also make CVL lists unique for particular Document Classes. For example, suppose you have two Document Classes: Programming Projects and Building Maintenance. You might want to set up both Document Classes to contain a custom property called Engineering Specifications. However, because two different sets of users will utilize the two Document Classes, you can make two different CVLs for the Engineering Specifications custom property, one for the Programming Projects Document Class and one for the Building Maintenance Document Class. This arrangement gives you a way to provide different sets of value selections appropriate for the two different kinds of users.

Notes

- (FNDTS00069901) Although you can create a CVL whose name contains a "<" or ">" character, Replication Services will fail when attempting to replicate the CVL. Don't use "<" or ">" in a CVL name if you need to replicate your CVLs.
- The DocType & CVL tool cannot import more than 10,000 values at a time into a CVL. If your import file contains more than 10,000 values, DocType & CVL Tool will import the first 10,000 values and then display a message to warn you that subsequent values in the file will be ignored.

The CS Document Class and CVL Admin is available via the CS Explorer Tools menu. Refer to the *CS Document Class and CVL Admin Online Help* for usage details.

System Messages

To ensure the functional efficiency of a library system and its associated user interfaces, such as CS Explorer, you must pay special attention to system-generated messages. The Messages property in the Administrators Group object is used to log error messages that are of particular interest to administrators. These logged messages provide a record of configuration discrepancies, disk-space and network-communications problems, and related system information.

To monitor what is happening in your library system, you should check these messages regularly. While some of the messages are primarily informational in nature, others will require your immediate attention to ensure that the system continues to function properly. Each Message contains properties such as From, Date, and Type, enabling you to easily determine which messages need immediate attention.

One type of message notifies you that an operation has completed. For example, when you retrieve an archived version file and place it back online for a user who has initiated a reclaim request for it, the system issues a message indicating that the operation has completed.

Another common type of message notifies you that an operation could not be completed. For example, suppose a user attempts to check in a version and the storage repository for that document's storage category is full. Not only does the library system notify the user at the time the problem occurs, it also sends a message to the Administrators group. You would probably respond immediately to this type of message by reconfiguring network resources. Then you could create one or more new storage categories and add them to the library system.

A third type of message informs you when the property repository or database transaction log reaches configurable capacity thresholds. For example, you may receive a warning message telling you that the property repository is now 80% full and you should assign more storage space to it as soon as possible to ensure that it continues to function normally. Refer to [“Configuration And Maintenance” on page 224](#) for details on how to respond to such threshold messages.

Only asynchronous system messages are displayed in the Messages property. Other types of messages may be displayed directly to you and other users as you work (for example, in CS Explorer).

Note For tool-specific instructions on how to view the message log, refer to the *CS Explorer Online Help*.

Suspending a Storage Manager

Suspending Storage Manager operation makes the Storage Manager, and the storage repositories assigned to it, unavailable to general users while allowing the library system administrator to perform maintenance tasks such as backups. While the Storage Manager is suspended, you can use any commands controlled by the System object. For example, you can add storage repositories, archive repositories, storage categories, and archive categories.

Once you have completed the administrative tasks, you can resume Storage Manager operation to make the library system available to users again. Or, if you shut down and restart the server on which a Storage Manager has been suspended, the Storage Manager's status automatically changes from Suspended to Running.

It's important to keep in mind that when you suspend a Storage Manager, no message is sent to alert the users that are logged on to the library system. To avoid problems (such as not being able to complete file operations in progress) you should email users to alert them of the upcoming suspension.

Note Procedures for suspending and resuming Storage Manager operation are provided in *CS Explorer Online Help*.

Utilities

The command-line utilities included with Content Services are listed below and are documented in the following sections.

Note The pathname documented for each utility goes through a directory named **filenet**, which is created when you run a new install of Content Services 5.5. If you upgraded to Content Services 5.5 from an earlier version, then this directory maybe named **idmds**.

| | |
|-------------------------|--|
| “ARCHDATE” on page 371 | Archives versions that were last accessed before a specified date. |
| “CAT_INDX” on page 375 | Repairs indexes in the database tables for the property repository. |
| “CAT_PROC” on page 377 | Deletes and redefines stored procedures in a Property Manager for a particular property repository. |
| “CAT_STAT” on page 379 | Updates usage statistics in the database engine for a particular property repository. CAT_STAT is available only for the Microsoft SQL Server edition of CS. |
| “CCAT” on page 380 | Loads data files containing scripts that update your property repository database tables. |
| “COPY_CP” on page 382 | Copies one or more custom properties from one library system to another. |
| “CSADMIN” on page 384 | Rebuilds, starts, and stops content search repositories. |
| “CSK2CFG” on page 386 | CSK2CFG enables you to repair a Verity configuration for a library system. |
| “CS_IMPORT” on page 392 | Parses an XML file, generated by the LDAP_EXP utility, into files of users and groups to be synchronized with those in a Content Services library system. |
| “DMU” on page 394 | Transfers items between storage repositories. |
| “DSINFO” on page 398 | Displays the status of library system services. |
| “DSSTOP” on page 400 | Stops library system services. |
| “DSSTRT” on page 402 | Starts library system services. |
| “DTCVLCHK” on page 405 | Detects Document Class or CVL (Controlled Vocabulary Lists) collisions between two or more library systems. |
| “DTCVLMIG” on page 412 | Migrates Document Classes and CVLs between two or more library systems. |
| “FILE2ID” on page 419 | Generates the associated item ID and version ID for the version file. |

| | |
|------------------------|--|
| “HMSRVCHK” on page 420 | Generates a list of subfolders and documents that will be sourced on Home Server on running HMSRVMIG and the documents and subfolders that have exceptions. |
| “HMSRVMIG” on page 424 | Sources all the documents and subfolders within a replication folder from member libraries to Home Server. |
| “ID2FILE” on page 428 | Generates the encrypted file name for a version. |
| “INDXDATE” on page 429 | Indexes for content search versions last accessed before a specified date. |
| “INDXFIX” on page 432 | Searches for and queues for indexing all versions with an indexing status of “filter failed” or “index failed”. |
| “LDAP_EXP” on page 435 | Binds to an LDAP directory, searches for requested group names, writes an XML file of user/group information to be parsed by the CS_IMPORT utility, and closes the LDAP directory. Add a display language for a library system. |
| “REPCHECK” on page 439 | Reports on or fixes links between the property repository and the storage or content search repositories. |
| “REPLKICK” on page 443 | Forces users, groups, Document Classes, and CVLs to replicate throughout a replication domain. It is run against the domain master and is used to synchronize library systems. |
| “SPISTAT” on page 445 | Prints the program-defined value for a specific error code. |
| “SRDELETE” on page 446 | Deletes items found by a given search. |
| “USGRPCHK” on page 448 | Detects User or Group collisions between two or more library systems. |
| “USGRPMIG” on page 455 | Migrates Users and Groups between two or more library systems. |
| “VUPDATE” on page 461 | Updates version file names. |

ARCHDATE

ARCHDATE enables you to archive versions that were last accessed (that is, checked out, checked in, copied, or viewed) before a specified date.

Usage

```
archdate /username= /password= /propertyserver= /systemname=  
/date= /itemsearch= /keepmainline= /forcearchcat=  
/substitutearchcat= /listonly= /logfile= /inputfile=
```

Parameters and Switches

| | |
|------------------|---|
| /username= | The user name of the person archiving versions. Only versions to which this user has at least Owner access rights can be moved to archive repositories. |
| /password= | The password, if any, for the user identified with the username parameter. If you use a batch file, we suggest that you not include the password parameter in your script (it can be read by anyone who can view the contents of that file). Instead, you should allow the utility to prompt you for this value (your response will be masked on the screen as you type it). |
| /propertyserver= | The name of the server where the database engine is installed. |
| /systemname= | The name of the CS library system. |
| /date= | The identifying archiving date. Versions that were last accessed before 12:00 a.m. of this date will be flagged for possible archiving. This date must appear in the format mm/dd/yyyy (for example, 01/21/1990). |
| /itemsearch= | (Optional) The name of a pre-existing item search. Note that ARCHDATE will produce inaccurate results if any search type other than an item search is specified for this parameter. Only versions of items returned by this search may be archived (based on the /date= parameter). |
| /keepmainline= | <ul style="list-style-type: none">• Yes. All versions last accessed before a specified date (except the current mainline version) are archived.• No. All versions last accessed before a specified date (including the current mainline version) are archived. This behavior matches the behavior in previous releases. |
| /forcearchcat= | <ul style="list-style-type: none">• Yes. The system archives all versions to the substitute archive category (set with the substitutearchcat parameter). Remember, the substitute archive category must always be one that is assigned to a <i>reclaimable</i> archive repository.• No. The system archives versions to the substitute archive category (set with the substitutearchcat parameter) only if the Default Archive Category in the Item object is assigned to a <i>nonreclaimable</i> archive repository. This behavior matches the behavior in previous releases. |

| | |
|---------------------|---|
| /substitutearchcat= | The substitute archive category that ARCHDATE will use to archive versions <i>only</i> if the default archive category in the version's parent Item object is linked to a nonreclaimable archive repository. This substitute archive category must be associated with a <i>reclaimable</i> archive repository. This parameter is required to ensure that archived versions are immediately reclaimable. |
| /listonly= | A Yes/No value to indicate whether you want to run the utility in list mode only (the default value is Yes). |
| /logfile= | The name of the file to which you want the log of the ARCHDATE results sent. This log will identify the item ID and version ID of the versions that were removed, the last access date of each of these versions, the number of bytes freed per version, and the archive category used. You can specify a value for this parameter in one of three ways: <ul style="list-style-type: none"> • Accept the default value, archdate.log. (If this file already exists, the new information will be appended to it.) • Specify a different log file name. (If the file already exists, the new information will be appended to it.) • Use con as the parameter value to indicate that you want the log printed to your monitor rather than to a file (for example, logfile=con). |
| /inputfile= | The name of the file containing prompt-response values for the command line parameters. |

Usage Notes

- On the command line, you must include a forward slash (/) before each parameter and a blank space between parameter statements. No space should appear between the parameter and its value; however, the parameter values themselves may contain blank spaces.

For example, you might type (on one line):

```
archdate /username=John Braughn
/propertyserver=01 /systemname=MZ1 /date=06/11/1995
```

and press Enter to start the utility. If you leave any parameters out (as in this example), you will be prompted for them individually.

- The versions to be archived will be identified according to a last-access date that you specify. This date will be compared to the Last Access Date property value in each Version object.
- The stored search request specified by ARCHDATE using the itemsearch parameter must be saved at the system level.
- If you do not use the itemsearch parameter when running ARCHDATE, all Item versions in the database that were last accessed before the specified date will be archived.
- The stored search request specified by ARCHDATE must have the “order by item id” clause set. Note that the first property from your search is the default sort order setting. (An exception to this is the keyword property which cannot be used to sort search results).

- The ARCHDATE utility also allows you to:
 - a. Archive versions based on a specified search as well as a last access date.
 - b. Run in list mode, enabling you to verify the results of a potential archive operation before you actually run it.
- To ensure that the archived versions are immediately *reclaimable*, this utility will archive versions to reclaimable archive repositories only. After the system determines that the user running the utility has at least Owner access rights to a version and that the version was last accessed before the date specified, the utility will archive the version in one of the following ways:
 - a. If the default archive category in the version's parent Item object is linked to a reclaimable archive repository, then the version is sent to that archive repository.
 - b. Otherwise, the version will be sent to a reclaimable archive repository that is linked to a substitute archive category (as specified by the required `substitutearchcat` parameter).

Because no version is ever sent to a nonreclaimable archive repository, there is less likelihood that you will inadvertently destroy a version.
- You can run ARCHDATE using either of the following:
 - a. Executing the utility from the command line.
 - b. Create a batch file and a corresponding prompt-response file from which you run ARCHDATE. This method allows you to start up the batch file (during low-usage times, for example) and let ARCHDATE archive versions based on the values you specify in the prompt-response file.

Procedures

To run ARCHDATE from the command line:

1. Review the ARCHDATE *Usage Notes* topic.
2. If you want to use a search to identify the versions to archive, you must create that search (and save it at the System access level) before running the ARCHDATE utility. Use the CS Admin Tools or any other interface to the library system.
3. On the initial storage server, access the directory containing the `archdate` executable.

| OS Platform | Location |
|-------------|---|
| Windows | <code>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</code> |
| UNIX | <code>StorageManagerInstallDir/filenet/librarysystem/util</code> |

4. On the command line, enter:

```
archdate
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

The search, if any, is run first, followed by the archival operations. The results are displayed or printed to the log file.

To run ARCHDATE from a batch file:

1. Review the ARCHDATE *Comments* topic.
2. If you want to use a search to identify the versions to archive, you must create that search (and save it at the System access level) before running ARCHDATE. Use the CS Admin Tools or any other interface to the library system.
3. Create a batch file with the following line in it:

```
archdate /inputfile=archdate.res
```

We recommend against including your library system password in the batch file. (Instead, allow the system to prompt you for it.) However, if you prefer to include the password so that this operation can be run completely unattended, replace the preceding batch file command with:

```
archdate /password=value /inputfile=archdate.res
```

If you do not have a library system password set, omit the password value:

```
archdate /password= /inputfile=archdate.res
```

4. Create the **archdate.res** file. In this file, type in your values for the command line parameters, excluding password and inputfile, in the correct order as shown in the following example. Include a space before and after equal signs (=) and separate each command line parameter with a carriage return.

```
username = John Doe
propertyserver = server01
systemname = docsrvr
date = 12/31/1995
forcearchcat = no
substitutearchcat = long term
itemsearch = archdate search
keepmainline = yes
listonly = no
logfile = archdat1.log
```

5. When you are ready to run ARCHDATE, start the batch file.

The search, if any, runs first. Then, as the archival operations take place, the results are either displayed on your monitor or printed to the specified log file. When all the identified versions have been moved to their appropriate archive repositories, you are returned to the command prompt.

CAT_INDX

CAT_INDX repairs corruptions that affect the indexes in the database tables for the property repository. With this utility, you can rebuild one index, the indexes for a particular table, or the indexes for an entire property repository.

Note CAT_INDX verifies that at least 30% of the total system tablespace is free before it starts any index rebuilding. If less than 30% of the total system tablespace is free, CAT_INDX issues a message recommending that you increase the tablespace size.

Usage

```
cat_indx propertyserver librarysystem tablename
```

Parameters and Switches

| | |
|-----------------------|--|
| <i>propertyserver</i> | The name of the server where the database engine is installed. |
| <i>librarysystem</i> | The name of the library system. |
| <i>tablename</i> | The database table (or database object ID) that identifies which indexes must be rebuilt. You can specify the indexes you want to rebuild by entering one of the values shown in the following list: all Rebuild indexes in all tables. <i>tablename</i> Rebuild indexes for this table. (Note: A database object ID may be used instead of the table name.) <i>tablename.index</i> Rebuild the specified index in the specified table. If you specify a <i>tablename.index</i> and the index is one that would exist in the schema but for some reason does not exist in your database, CAT_INDEX will re-create it. Note An index ID may be used instead of the index name. This ID is published in the errorlog file as INDID=x. |

Usage Notes

- Because this utility can potentially affect the future availability of your stored information, we recommend you use it only after thoroughly reading through this documentation or checking with FileNet Technical Support.
- Typically, CAT_INDX is used to rebuild indexes that have row-length problems.

Procedures

To prepare to run CAT_INDX:

1. Notify all users to close their library system sessions.
2. Make sure the Logons Enabled? property in the System object is set to No.

This allows only one user from the Administrators group to open a library system session.
3. Run and process a full system backup (selecting the Consistency Check option to allow the CAT_FIX utility to run).

To run CAT_INDx:

1. On the initial storage server, access the directory containing the **cat_indx** executable.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

2. On the command line, enter:

```
cat_indx
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

As the utility runs, it will display a message indicating the index that is being built for the specified table. If an error occurs, an appropriate message will be displayed.

3. When the utility finishes, set the value of the Logons Enabled? property in the System object to Yes. This allows other users to access the library system.

Examples

- If you wanted to rebuild the index I12 in the Element (that is, item) table for the system Lib1 on the server Sys1, you would type the following command on the command line:

```
cat_indx sys1 lib1 element.i12
```

- If you had only the database object ID (for example, 97849382) instead of the table name and the index ID (for example, 14) instead of the index name, you could type:

```
cat_indx sys1 lib1 97849382.14
```


CAT_PROC

CAT_PROC enables you to delete and redefine the stored procedures (that is, precompiled operations) for a particular property repository.

Usage

```
cat_proc propertyserver librarysystem
```

Parameters and Switches

| | |
|-----------------------|--|
| <i>propertyserver</i> | The name of the server where the database engine is installed. |
| <i>librarysystem</i> | The name of the library system. |

Usage Notes

- Because this utility can potentially affect the future availability of your stored information, we recommend you use it only after thoroughly reading through this documentation or checking with FileNet Technical Support.
- You can run CAT_PROC only on a initial storage server. If you run CAT_PROC on another library system server, it will not find the necessary configuration file and will display the error: “Failed to read dsconf.ini.”
- CAT_PROC can be useful in those rare cases when one or more stored procedures in the property repository are damaged. Such damage may be indicated when a particular library system operation, without explanation, fails to complete successfully. For example, if the stored procedure for checking in versions were damaged, then each time the user attempted this operation, he or she would typically receive a general internal system error.
- CAT_PROC uses an internal source file to reload the stored procedures. You must have at least 4000 kilobytes of free space in the transaction log to run this utility.

Procedures

To prepare to run CAT_PROC:

1. Notify all users to close their library system sessions.
2. Make sure the Logons Enabled? property in the System object is set to No.
This allows only one user from the Administrators group to open a library system session.

To run CAT_PROC:

1. On the initial storage server, access the directory containing the **cat_proc** executable.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

2. On the command line, enter:

```
cat_proc
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

3. When the utility finishes, set the value of the Logons Enabled? property in the System object to Yes. This allows other users to access the library system.

CAT_STAT

CAT_STAT enables you to update usage statistics for a particular property repository.

Note CAT_STAT is available only for the Microsoft SQL Server edition of CS.

Usage

```
cat_stat propertyserver librarysystem
```

Parameters and Switches

| | |
|----------------|--|
| propertyserver | The name of the server where the database engine is installed. |
| librarysystem | The name of the library system. |

Usage Notes

- Usage statistics that CAT_STAT updates can affect the accuracy of the system's estimates of warning and full thresholds.
- You can run CAT_STAT manually from the command line, or you can allow the system to run it automatically during full system backups.
- Typically, you will allow the system to run this utility as part of a full backup. This automatic running of the utility is usually adequate to keep property repository usage statistics up to date.

Procedure

To run CAT_STAT:

1. Review the CAT_STAT *Usage Notes* topic.
2. On the initial storage server, access the directory containing the **cat_stat** executable, located in:

```
StorageManagerInstallDrive:\filenet\LibrarySystem\Util
```

3. On the command line, enter:

```
cat_stat
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

CCAT

CCAT loads data files (*.dat) containing scripts that update your system's database tables.

Note CCAT is not a general purpose tool. It should generally be used only with the assistance of FileNet Technical Support.

Usage

```
ccat -U libraryname -D libraryname -S propertyserver -q -x -z -p -n -B
datfilename -X logfilefilename -L 16
```

Parameters and Switches

Procedure

| parameter | meaning |
|-----------------------------|--|
| -U <i>libraryname</i> | DB User name, the account created while running loginscript.sql when library is installed, it is same as the library name. |
| -D <i>libraryname</i> | DB name, it is same as library name. |
| -S <i>propertyserver</i> | Machine name on which database is created. |
| -P | DB Password [optional] |
| -q | Quit on Error [optional] |
| -x | Encryption on (Required for dat file) |
| -z | Screen on [optional] |
| -p | No Prompt [optional] |
| -n | No SP check [optional] |
| -B | Batch File Name to be executed |
| -X | Log file name [optional] |
| -L | Logging ON [optional] with bit options |

To run CCAT:

1. Inform all users to close their library system sessions.
2. On the initial storage server, access the directory containing the **ccat** executable.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

3. On the command line, enter:

```
ccat -s propertyserver -U librarysystem -D librarysystem -x filename.dat
```

4. A series of SQL statements will scroll on the screen. If no error codes are returned, the script file has been loaded successfully.
 - a. Repeat steps 1 through 4 for any other library systems whose database tables need to be updated.
 - b. When you are done executing CCAT, set the value of L_LOGINS_ENABLED in LIBRARY Table (SPI_OBJ_SYSTEM, 1280) to YES. This allows other users to access the Library System.

COPY_CP

COPY_CP enables you to copy one or more custom properties from one library system to another, thereby ensuring that documents and stored searches are replicated properly between library systems.

Caution Once you add a custom property to any library system, it cannot be removed.

Note COPY_CP is currently available only for Windows, however, the Windows version can be used with UNIX library systems. COPY_CP.exe is available on the *FileNet Content Services Admin Tools* CD. If you are running Replication on UNIX servers and do not have at least one Windows server to run this utility, you will have to manually create/duplicate any necessary custom properties on the various library systems participating in replication services.

Usage

copy_cp

Usage Notes

- You may create custom properties (instances of the Custom Property object) as document properties, which users can then use to identify, or create searches to find, the documents in the library system. The *value* of a custom property will be replicated with the associated document or stored search, but only if you have already created a custom property of the same name and type on the target library system. FileNet Replication will not replicate custom properties automatically, since they may not exist, or may have different values on other library systems participating in replication services.
- The custom property utility copies one or more custom properties from a source library system to a destination library system. It will not overwrite any matching custom properties on a library system. Additionally, once the maximum of number of custom properties allowed per library system is reached (refer to “[Custom Properties](#)” on page 258), the utility will copy no further custom properties.
- There are no parameters for the custom property utility.

Prerequisites

- For any library system to which you intend to copy custom properties, you must perform a full backup operation before running the custom property utility.
- You must run the custom property utility *before* enabling replication services for the library system in question.
- You must run the custom property utility from the FileNet Replication installation directory.
- You must have Administrators access rights to both the source and destination library systems specified for the custom property utility.

Procedure

To copy custom properties:

1. Do one of the following:
 - (Windows) On the initial storage server, access the directory containing the **copy_cp.exe** utility:

Drive:\filenet\Repl\Util

- (Unix) On a Windows NT administrative workstation, access the directory containing the **copy_cp.exe** utility on the *Content Services Admin Tools* CD:

Drive:\Util

2. Start the copy_cp utility.
3. Answer the following prompts:
 - Source/Destination server name
 - Source/Destination system name
 - Source/Destination Administrators user name
 - Source/Destination Administrators password

4. When prompted, specify:

| | |
|------|--|
| Yes | Confirm each copy |
| No | Proceed without confirmation for the displayed property only |
| All | Copy all custom properties without confirmation |
| Quit | Exit without copying |

CSADMIN

CSADMIN rebuilds, starts, and stops content search repositories.

Caution For Microsoft SQL Server databases, running CSADMIN in rebuild mode has the potential to fill up your transaction log. One way to avoid this problem is to enable the “Truncate log on checkpoint” feature before running CSADMIN (and then disable it afterward if desired). This option is available in Enterprise Manager, on the Options tab of the database properties sheet. Another way to avoid the problem is to configure the transaction log file to auto-grow. This option is available in Enterprise Manager, on the Transaction Log tab of the database properties sheet. If you use this method, you should monitor the process to ensure that the transaction log file does not fill up your disk. Refer to *SQL Server Books Online* for more information about these database features.

Usage

```
csadmin propertyserver librarysystem -i -r -s -t -v -a
```

Parameters and Switches

| | |
|-----------------------|---|
| <i>propertyserver</i> | The name of the server where the database engine is installed. |
| <i>librarysystem</i> | The name of the library system. |
| -i | (Optional) Resets the status of the indexing daemon to idle. It will then start at the next scheduled indexing interval. |
| -r | (Optional) Rebuilds the content search repositories. You are prompted how to rebuild the content search repositories: <ul style="list-style-type: none">• From a full backup A consistency check is performed on the content search repository between the last full backup and the current content search repositories.• From the last fully restored backup of repository files.• From scratch. |
| -s | (Optional) Suspends the indexing daemon after the indexing operation being done on the current version file has completed. If a series of files were being processed, the first file in the series would be completed before the repository would shut down. |
| -t | (Optional) Terminates content search indexing. |
| -v | (Optional) Verifies that the content search server contains indexes. |
| -a | Rebuilds all repositories from scratch (Without prompting for each). |

Usage Notes

- We recommend you use CSADMIN only after thoroughly reading through this documentation or checking with FileNet Technical Support.
- Command line switches are all case insensitive.

- On average, indexing 1000 versions will consume 70 MB of transaction log space. We recommend setting your log full threshold to 80% or less and indexing versions in increments so that the log does not fill to capacity.
- After running CSADMIN to rebuild content search repositories on a Windows server, reboot the server.

Examples

To rebuild a content search repository for a library system named corporate on server corp1, you would enter the following command and then respond to the CSADMIN prompts.

```
csadmin corp1 corporate -r
```

Procedure

To run CSADMIN:

1. Review the CSADMIN *Usage Notes* topic.
2. On the content search server, access the directory containing **csadmin.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>CSMInstallDrive:\filenet\LibrarySystem\Content\Bin</i> |
| UNIX | <i>CSMInstallDir/filenet/librarysystem/content/bin</i> |

3. On the command line, enter:

```
csadmin propertyserver librarysystem switches
```

and respond to the CSADMIN prompts.

4. Reboot the server where you ran CSADMIN. If you are running CSADMIN more than once on the server, to rebuild more than one content search repository, you need only reboot once after the last execution of CSADMIN.

CSK2CFG

CSK2CFG enables you to verify and repair a Verity configuration for a library system.

The purpose of the CSK2CFG utility is to verify and repair a Verity component configuration, and to do so in a way that is simpler than using the Verity command-line tool `rcadmin`. CSK2CFG reads the CS database tables and the Windows registry or Unix configuration file to obtain information about the Verity configuration and to make repairs.

You can invoke this utility whenever you suspect that the Verity configuration has become defective, as indicated by indexing or content-search failures or when any aspect of the configuration fails during (re) install/upgrade of CS software. The following procedures involve the configuration of Verity components:

- Installing or upgrading an initial or additional set of services for a library system.
- Moving a set of initial or additional services for a library system from one server to another.
- Renaming a storage server.
- Moving the database from one server to another.
- Moving a Master Admin server from one Windows domain to another.
- Adding, moving, or restoring a content search repository.

Note A Master Admin server is a storage server where a Verity K2 Master Admin is installed.

Each of the procedures listed above, except the last one, involves running the CS installation software, which results in the creation of a log file (`cssetup.log`) with entries indicating whether the Verity configuration succeeded or failed. Under Windows, this log file is located in **InstallDrive:\filenet\cssetup.log**.

A proper configuration of Verity components installed with a set of services for a library system on a server is characterized as follows:

- A Verity K2 Broker and a Verity K2 Server are installed and registered with the local Verity K2 Admin on the server where the (re)install/upgrade takes place. (Registration means that the K2 Admin is aware of the existence of the K2 Broker and K2 Server).
- The Verity K2 Server is attached to the Verity K2 Broker.
- A Verity collection (content search repository file named `index0xx`) is created, registered with the Verity K2 Admin, and attached to the Verity K2 Server.

If any aspects of the Verity configuration failed during a (re)install/upgrade of CS software, as indicated in `cssetup.log`, you can repair the configuration by running the utility CSK2CFG as described in the following sections.

Note Occasionally the `cssetup.log` file will not indicate any Verity configuration problems when you install an additional set of library services. But indexing and/or searching in the library system fails. In such cases, you need to run CSK2CFG on the additional storage server to repair the configuration.

To obtain the complete status of a Verity configuration for a given library system without changing any configuration settings, invoke CSK2CFG, without specifying any parameter switches, on each server where initial or additional services of the library system are installed, as follows:

```
csk2cfg dbserver libname
```

where *dbserver* specifies the database (property) server for the library system *libname*.

Note How you specify the database server depends on the server operating system: If the operating system is Windows, then *dbserver* is the ODBC data source name. If it is UNIX (HP-UX or Sun Solaris), then *dbserver* is the Oracle net service name.

The output of CSK2CFG tells you whether the Verity components are properly configured for a particular set of initial or additional library services. Once you have determined that the configuration is faulty, you can run CSK2CFG with the attach-mode parameter *-A* to repair the Verity components for that set of library services:

```
csk2cfg dbserver libname -A
```

If you want to create a fresh Verity configuration for your library system, you will first need to run CSK2CFG with the detach-mode parameter *-D* on each server where services for the library system are installed:

```
csk2cfg dbserver libname -D
```

You can then run CSK2CFG with the attach-mode parameter *-A* on each server where services for the library system are installed.

Usage

```
csk2cfg dbserver libname
```

```
csk2cfg dbserver libname -A
```

```
csk2cfg dbserver libname -D
```

Parameters and Switches

| | |
|-----------------|--|
| <i>dbserver</i> | Identifies the server where the database engine is installed, based on the operating system running on the server. For Windows, the identifier is the ODBC data source. For UNIX, the identifier is the Oracle net service name. |
| <i>libname</i> | Identifies the library system whose Verity components are to be repaired. |
| <i>-A</i> | Registers and attaches Verity components of a library system. |
| <i>-D</i> | Un-registers and detaches Verity components of a library system. |

Usage Notes

- Run CSK2CFG with at most one switch per invocation.
- Switches are all case insensitive.
- When repairing the Verity components for a library system, you can invoke CSK2CFG for the affected sets of initial and additional services in any order.

Procedure

To run CSK2CFG:

1. Verify that the database services are running on the database server.

2. Verify that the local K2 Admin for each set of library services installed on the storage server is running. If it is not running, start it as follows:
 - (Windows storage servers) Start the CS Verity K2Admin service in the Administrative Tools > Services panel.
 - (Unix storage servers) Run k2adminstart, a shell script in the \$IDMDS_HOME directory.
3. Log on to the storage server of the library system on which you will be running CSK2CFG.
4. Access the directory containing CSK2CFG.

| OS Platform | Location |
|-------------|---|
| Windows | <i>CSMInstallDrive:\filenet\libname\content\bin</i> |
| UNIX | <i>CSMInstallDir/filenet/libname/content/bin</i> |

5. On the command line, type:

```
csk2cfg
```

and the desired parameters and switch.

Examples

The following examples assume a newly installed or upgraded library system lib1 using database server svr0, initial storage server svr1, and additional storage servers svr2, svr3, and svr4.

Example 1: Installing or upgrading a set of library services

After installing or upgrading a set of initial or additional services for library system lib1, you need to determine if the Verity components that are part of the installation/upgrade are correctly configured

1. On the storage server where initial or additional services for lib1 are installed, run CSK2CFG to determine the status of the Verity components, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1
```

2. If the Verity components need repair, run the CSK2CFG in attach mode on the server where the newly installed/upgraded services are located, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -A
```

Example 2: Moving a set of initial library services from one server to another

Suppose you have moved the set of initial services for library system lib1 from server svr1 to svr5. You must now run CSK2CFG in attach mode to restore a correct Verity configuration:

1. On each storage server where initial or additional services for lib1 are installed, run CSK2CFG in attach mode, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -A
```

2. If the set of initial services you moved to svr5 contains a Verity K2 Master Admin that governs library systems besides lib1, do the following:

- a. Identify the initial and additional storage servers for each of the other library systems governed by the Verity K2 Master Admin you moved to svr5.
- b. On each storage server identified in step a, run CSK2CFG once for each set of initial and additional library services on the server, invoking CSK2CFG as follows:

```
csk2cfg srv0 libname -A
```

where *libname* is the name of a library system whose initial or additional services are on the storage server where CSK2CFG is run.

Example 3: Moving a set of additional library services from one server to another

Suppose you have moved a set of additional services for library system lib1 from server svr4 to svr5. You must now run CSK2CFG in attach mode to restore a correct Verity configuration:

1. On storage server svr5, run CSK2CFG in attach mode, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -A
```

Example 4: Renaming a storage server

Suppose you have renamed a storage server. You must now run CSK2CFG in attach mode to restore a correct Verity configuration:

1. Identify the sets of initial and additional library services installed on the renamed storage server. If the storage server only has sets of additional services installed on it, continue at step 2 and skip the remaining steps. If it has at least one set of initial services installed on it, continue at step 3.
2. Run CSK2CFG once for each set of additional library services on the renamed server, invoking CSK2CFG as follows:

```
csk2cfg srv0 libname -A
```

where *libname* is the name of a library system whose additional services are on the renamed server.

3. Run CSK2CFG once for each set of initial and additional library services on the renamed server, invoking CSK2CFG as follows:

```
csk2cfg srv0 libname -A
```

where *libname* is the name of a library system that has initial or additional services installed on the renamed server.

4. Identify the storage servers where there are additional library services corresponding to the initial library services on the renamed server.
5. On each storage server identified in step 4, run CSK2CFG once for each set of additional library services corresponding to a set of initial library services on the renamed server, invoking CSK2CFG as follows:

```
csk2cfg srv0 libname -A
```

where *libname* is the name of a library system whose initial services are on the renamed server.

If a Verity K2 Master Admin is on the renamed storage server, and it governs library systems whose initial services are on other storage servers, then continue at step 6; otherwise, you are done.

6. Identify the initial and additional storage servers for each of the other library systems governed by the Verity K2 Master Admin on the renamed storage server.
7. On each storage server identified in step 6, run CSK2CFG once for each set of initial and additional library services on the server, invoking CSK2CFG as follows:

```
csk2cfg srv0 libname -A
```

where *libname* is the name of a library system whose initial or additional services are on the storage server where CSK2CFG is run.

Example 5: Moving the database to another server

After moving the database from server *svr0* to server *svr5*, you must run CSK2CFG in attach mode on each storage server where there are library services supported by the database. Thus, for each library system *libname* supported by the database, do the following:

1. Identify the storage servers where there are sets of initial or additional services of library systems supported by the database.
2. On each storage server identified in step 1, run CSK2CFG in attach mode once for each set of initial and additional library services installed on the server, invoking CSK2CFG as follows:

```
csk2cfg srv5 libname -A
```

where *libname* is the name of a library system whose initial or additional services are on the storage server where CSK2CFG is run.

Example 6: Moving a Master Admin server from one Windows domain to another

Suppose that you have moved the initial storage server *svr1* for library system *lib1* from one Windows domain to another, and that a Verity K2 Master Admin is part of the initial services of *lib1*. You must now run CSK2CFG in attach mode to restore a correct Verity configuration for the initial and all additional services of *lib1* and of every other library system governed by this Verity K2 Master Admin:

1. Identify the initial and additional storage servers for each library system governed by the Verity K2 Master Admin on *svr1*.
2. On each storage server identified in step 1, run CSK2CFG once for each set of initial and additional library services on the server, invoking CSK2CFG as follows:

```
csk2cfg srv0 libname -A
```

where *libname* is the name of a library system whose initial or additional services are on the storage server where CSK2CFG is run.

Example 7: Adding, moving, or restoring a content search repository

After adding, moving, or restoring a content search repository to a storage server, you must run CSK2CFG on the server to register and attach the repository (Verity collection) to the Verity K2 Server on the server:

1. Run CSK2CFG in attach mode on storage server where you added the repository, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -A
```

Example 8: Verifying and repairing an existing installation

If you suspect that Verity indexing or searching is no longer working on library system lib1, you can diagnose the problem by running CSK2CFG to determine the status of the Verity components configured for each set of services for the library system.

1. On each storage server where initial or additional services for lib1 are installed, run CSK2CFG to determine the status of the Verity components configured for the lib1 services, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1
```

2. For each set of services for lib1 whose Verity components need repair, run the CSK2CFG in attach mode on the server where the services are located, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -A
```

Example 9: Detaching and re-attaching a Verity configuration

Suppose you have noticed that Verity indexing or searching is no longer working on library system lib1. Rather than repairing the Verity configuration on one or two storage servers for lib1, you choose to repair the Verity configuration by generating a fresh configuration for all the storage servers.

1. On each storage server where initial or additional services for lib1 are installed, run CSK2CFG in detach mode, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -D
```

2. On each storage server where initial or additional services for lib1 are installed, run CSK2CFG in attach mode, invoking CSK2CFG as follows:

```
csk2cfg srv0 lib1 -A
```

CS_IMPORT

Parses an XML file, generated by the LDAP_EXP utility, into files of users and groups to be synchronized with those in a Content Services library system. Refer to “Importing Users and Groups” on page 272 for information on using CS_IMPORT and LDAP_EXP.

CS_IMPORT also provides an option for importing users/groups using multiple XML files that are created using IS LDIF tool: fn_ldif_xfer. Refer to “XML User Import using LDIF (LDAP Data Interchange Format) Format Files” on page 289 for more information.

Usage

```
cs_import /h server:library /u user /p password  
/i XML-input-filename /t /? /b
```

Parameters and Switches

| | |
|-----------------------|---|
| <i>server:system</i> | The name of the server and library system whose users/groups are to be updated by the users/groups specified in the XML file. |
| <i>user</i> | Administrative user who will log in to the library system and update its users/groups. |
| <i>password</i> | Encrypted password of the administrative user. Must be delimited by double-quotes. |
| <i>XML-input-file</i> | Pathname of XML file containing list of users/groups to import. |
| <i>t</i> | Generates a preview file (not enabled) |
| <i>?</i> | Generates a help window for the command-line arguments. |
| <i>b</i> | Batch-mode operation to import users/groups from multiple XML files that are created using IS LDIF tool: fn_ldif_xfer. |

Usage Notes

- CS_IMPORT is strictly a command-line oriented program; and all its required parameters are passed to it via the command line.
- You can use services of the operating system to automatically run CS_IMPORT at specified intervals. In Windows, use the Schedule service. In Unix, use the cron command.
- Each parameter and switch is required unless otherwise noted.
- The command switches may be prefixed by either '-' or '/'.
- To obtain your encrypted password, invoke the command-line utility pwcrypt, located in the filenet\systemname\util directory, as follows:

```
pwcrypt plaintext_password
```

The output of the pwcrypt command is

```
"plaintext_password" password is "encrypted_password" encrypted.
```

Example:


```
pwcrypt foobar
```

"foobar" password is "<,=8//" encrypted.

- The XML input file will be renamed to **cs_yyyymmddhhmmss.xml** after each run for tracking purposes.
- CS_IMPORT does not import LDAP subgroups or their members. That is, when you specify an LDAP group (in the LDAP User Import tab) to be imported/synchronized with the groups and users in a library system, only the LDAP group and its immediate members (users) can be imported, not the subgroups or their members.
- A standard log file, **LDAP_import_yyyymmddhhmmss.log**, will be created or appended to for each run.
- Within the current directory, three ASCII sort files will be created during each run:
 - **LDAP_import_yyyymmddhhmmss.usr** contains all the unique users in the run.
 - **LDAP_import_yyyymmddhhmmss.grp** contains all the unique groups in the run.
 - **LDAP_import_yyyymmddhhmmss.mem** contains all the unique user-in-group memberships in the run.

Procedure

To run CS_IMPORT:

1. Review the CS_IMPORT *Usage Notes* topic.
2. Access the directory containing CS_IMPORT.

| OS Platform | Location |
|-------------|--|
| Windows | <i>CSMInstallDrive:\filenet\LibrarySystem\util</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/util</i> |

3. On the command line, enter:

```
cs_import -h server:system -i XML_file_name -u AdminUser -u password
```

DMU

DMU enables you to move versions of items from one storage repository to another, using a transfer repository as an intermediate archive. DMU is a command-line utility. You can run it on a library system administrative client or any machine that has the CS Client Library Files installed and configured to connect to the affected library system. Before you can move versions to a target storage repository, you must already have archived them to a transfer repository (see “Moving Versions Between Storage Repositories” on page 236 for details.)

Usage

You can run DMU in four modes, depending on what you want it to do.

1. The following command lists the IDs of the archive repositories in a library system, so you can see which one will serve as the transfer repository.

```
dmu -A propertyserver librarysystem logfile
```

2. The following command lists the IDs of the storage repositories in a library system, so you can decide which one will serve as the target repository.

```
dmu -S propertyserver librarysystem logfile
```

3. The following command lists the item and version IDs of the versions you have archived to the transfer repository, without actually moving them to the target repository. Use this command to verify that you have correctly archived all the versions you wish to move to the target repository.

```
dmu -t propertyserver librarysystem logfile transfer_repository ID  
target_repository_ID
```

4. The following command moves the versions from the transfer repository to the target storage repository.

```
dmu -x propertyserver librarysystem logfile transfer_repository ID  
target_repository_ID admin_user admin_pwd
```

Note The admin user and admin password parameters must be used as a pair; neither can be used on the command line without the other. If you do not specify them on the command line, DMU will terminate with a message that displays the command syntax.

Note A *frozen* storage repository should not be used as the target storage repository.

Parameters and Switches

| | |
|-------------------|---|
| -A | List the archive repositories in a library system. |
| -S | List the storage repositories in a library system. |
| -t | List the versions that will be moved from the transfer repository to the target storage repository, without actually moving them. |
| -x | Move the versions from the transfer repository to the target storage repository. |
| <i>logfile</i> | A text file that contains a report of the DMU command execution. |
| <i>admin_user</i> | The user name for the library system Admin account |

| | |
|------------------------------------|---|
| <i>admin_pwd</i> | The password for the library system Admin account |
| <i>librarysystem</i> | The library system that contains the storage repositories involved in the movement of versions. |
| <i>propertyserver</i> | The property server that contains the library system. |
| <i>transfer_repositor y_ID</i> | The ID of the archive repository that serves as the transfer repository. |
| <i>target_repository_ ID</i> | The ID of the target storage repository. |

DMU Messages

Note In the following error messages %d and %ld are decimal numbers. They represent a specific repository ID, a count of version, or (most often) an error code, depending on the context. On the other hand, %s.%s indicates an item ID followed by a version ID.

Login errors:

- Username or password is incorrect.

DMU could not log on to the library system because the administrative user name or password was incorrect. Run DMU again with the correct user name and password.

- Error %d during login.

DMU could not log on to the library system. Possible causes are that the Property Manager isn't running, DMU can't connect to the Property Manager, or the CS Client Library Files or database client files on the DMU machine are either corrupt or the wrong version. Ensure proper database connectivity and run DMU again.

Errors in designating the transfer repository:

- Bad archive repository ID %d.

The archive repository ID given on the command line for the transfer repository is invalid. Try running `dmu -A` to see a list of valid IDs. Determine the correct ID and run DMU again.

- Error %d validating archive repository ID.

There was a failure while checking the transfer repository ID. Ensure proper database connectivity and run DMU again. If the error occurs again, it may be because the library system is damaged. You may need the help of FileNet Customer Support Services to recover from this situation.

Errors in designating the target storage repository:

- Bad storage repository ID %d.

The storage repository ID given on the command line for the target repository is invalid. Try running `dmu -S` to see a list of valid IDs. Determine the correct ID and run DMU again.

- Error %d validating storage repository ID.

There was a failure while checking the target repository ID.

Errors with versions in the transfer repository:

- Error %d getting count of versions in the designated archive repository.

There was a failure while retrieving the number of versions in the transfer repository.

- Nothing to do!

No versions were found in the transfer repository, so there's nothing to do to.

- Error %d updating storage repository IDs.

There was a failure while updating the storage repository IDs of the versions in the transfer repository.

Depending on the cause of the error, it may be possible to simply run DMU again with the same arguments to continue the process, after the problem is rectified. At worst, it should be possible to reclaim the versions back into their original storage repositories using the reclaim functionality provided by IDM Desktop.

- Consistency failure: originally got %ld versions in designated archive repository, but only updated %ld versions. Attempting to continue.

The number of versions whose storage repository IDs were updated does not match the number of versions in the transfer repository (which was retrieved earlier). This should never happen.

Reclaim errors:

- Error %d retrieving item and version ids.

There was a failure while setting up to reclaim the versions in the transfer repository. Depending on the cause of the error, it may be possible to simply run DMU again with the same arguments to continue the process, after the problem is rectified. At this point, the only thing that has been done is the updating of the storage repository IDs to point at the target repository, and there's no harm in doing that again.

- %s.%s: reclaim pending.

This file was removed from the archive repository and must be replaced. Once you replace it, the file will automatically be reclaimed. During the reclaim, the file associated with the version is not present in the transfer repository.

Note If this error appears, it will probably be necessary to restore the file not to the transfer repository but rather to some other archive repository. After DMU finishes, check the item to see what its archive category is, then check that category to identify its associated archive repository. It will probably be necessary to restore the file to that repository instead.

- %s.%s: reclaim FAILED, error %d

There was some other failure during the reclaim. Depending on the cause of the error, it may be possible to reclaim the file using the reclaim functionality provided by IDM Desktop.

Note The archive repository ID will not be reset! It may be necessary to reset it using a SQL query tool under the guidance of FileNet Customer Support Services.

- Consistency failure: originally got %ld versions in designated archive repository, but only found %ld to reclaim.

The number of versions reclaimed does not match the number of versions in the transfer repository (which was retrieved earlier). This can happen if users are active on the system and a user has caused the version to be reclaimed. Note that if another “consistency failure” message appears, it will precede this message. If the second number in both messages matches, then the problem is as described for the previous message. If no previous “consistency failure” appears, or the second number in this message is lower than the second number in the previous message, then some versions were reclaimed after their storage repositories were updated to the target repository, but before DMU could reclaim them and reset their archive repositories. The result is that the versions now are in the target repository as desired, but their archive repositories still point at the transfer repository. It may be necessary to update them with a SQL query tool under the guidance of FileNet Customer Support Services.

- Consistency failure: originally got %ld versions in designated archive repository, but only updated %ld versions. Attempting to continue.

The number of versions whose storage repository IDs were updated does not match the number of versions in the transfer repository (which was retrieved earlier). This can happen if users are active on the system and a user has caused the version to be reclaimed. In that event, the only thing that has occurred to the version is that its archive repository has been set to the transfer repository. At some time it can be archived again, using either ARCHDATE or the desktop (since its archive repository has been changed), and then DMU run again to move it to the target repository.

Archive repository reset errors:

- Consistency failure: there is more than one version with ID %s.%s?

While resetting the archive repository ID of the version shown, the database reported that more than one version was updated. Since the combination of item and version ID is supposed to be unique, this should never happen.

- Error %d resetting archive repository ID for %s.%s.

There was a failure while resetting the archive repository of the version shown. It may be necessary to reset it using a SQL query tool under the guidance of FileNet Customer Support Services.

Procedure

To run dmdu:

1. If you haven't done so already, copy the **dmdu.exe** file from the `\util` directory of the CS Admin Tools CD to a machine where the CS Client Library Files for Windows are installed (e.g., a CS Admin Tools client). You must also have configured a connection from this machine to the database server for each library containing the documents versions you will be moving.

Note On Windows editions of CS, you can find the **dmdu.exe** utility in the `\filenet\libraryname\util` folder.

2. Access the directory where the **dmdu.exe** file is located.
3. On the command line, enter:

```
dmdu
```

along with the appropriate switches and parameters.

DSINFO

Description

DSINFO displays the status of one or more components for one or more library systems on the server where it is run.

Usage

```
dsinfo -h
```

```
dsinfo -i librarysystem -c components -g NO
```

Flags and Parameters

| | |
|----------------------|---|
| -h | Causes help text to be displayed. Cannot be used with other flags. |
| -i | Must precede the <i>librarysystem</i> parameter. |
| <i>librarysystem</i> | The name of the library system. If you do not specify a name, dsinfo displays the status of the specified services for all library systems on the server where it is run. |
| -c | Must precede the <i>components</i> indicators. |
| <i>components</i> | You can include one or more of the following indicators to specify the components for which you want to see a status report. R Replication Service RA All Replication Services on the specified server S Storage Manager SA All Storage Managers on the specified server T Content Search Manager TA All Content Search Managers on the specified server |
| -g | When followed by NO, specifies no screen output. |

Usage Notes

- If you do not specify any component values, dsinfo provides status information for all services for the specified library system.
- If a service is not running, you may find that it is either shut down or suspended.
- Because flags are used to identify each parameter, you can type the parameters in any order, as long as the correct flags are included.

Procedure

To run dsinfo:

1. On the initial storage server or additional storage server, access the directory containing **dsinfo.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet</i> <i>StorageManagerInstallDrive:\Ifilenet\LibrarySystem\Util</i> (for backward compatibility) |
| UNIX | <i>\$IDMDS_HOME</i> <i>StorageManagerInstallDir/filenet/librarysystem/util</i> (for backward compatibility) |

2. On the command line, enter:

```
dsinfo
```

and the desired parameters and flags. If you do not enter any flags or parameters, you will not be prompted. Status information is displayed on your monitor.

Examples

Refer to “[DSSTRT](#)” on page 402 for examples of how to use dsinfo. The dsinfo, dsstop, and dsstrt utilities have identical parameters and flags.

DSSTOP

DSSTOP stops one or more services for one or more library systems on the server where it is run.

Note DSSTOP will remove all current active and inactive library system sessions. Keep this in mind when writing or working with custom applications (for example, if you plan to use DSSTOP in a backup script).

Usage

```
dsstop -h
dsstop -i librarysystem -c components -g NO
dsstop -i -d -g NO
```

Flags and Parameters

| | |
|----------------------|---|
| -h | Causes help text to be displayed. Cannot be used with other flags. |
| -i | Must precede the <i>librarysystem</i> parameter. |
| <i>librarysystem</i> | The name of the library system. If you do not specify a name, dsstop stops the specified services for all library systems on the server where it is run. |
| -c | Must precede the <i>components</i> indicators. |
| <i>components</i> | You can include one or more of the following indicators to specify services you want to stop. |
| | R Replication Service |
| | RA All Replication Services on the specified server |
| | S Storage Manager |
| | SA All Storage Managers on the specified server |
| | T Content Search Manager |
| | TA All Content Search Managers on the specified server |
| -d | Clears the sessionable and resets counters (for example, NUMID counter for sessions). May be used alone, or with “-i <i>library system</i> ” switch to clear individual library systems. The flag cannot be used with -c. |
| -g | When followed by NO, specifies no screen output. |

Usage Notes

- If you do not specify any component values, dsstop stops all services for the specified library system.
- Because flags are used to identify each parameter, you can type the parameters in any order, as long as the correct flags are included.

Procedures

To run dsstop:

1. On the property or storage server, access the directory containing **dsstop.exe**.

| OS Platform | Location |
|-------------|--|
| Windows | <i>StorageManagerInstallDrive:\filenet</i> <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Pbin</i> (for backward compatibility) |
| UNIX | <i>\$IDMDS_HOME</i> <i>StorageManagerInstallDir/filenet/librarysystem/pbin</i> (for backward compatibility) |

2. On the command line, enter:

```
dsstop
```

and the desired parameters and flags. You must enter the parameters and flags on the command line (that is, you will not be prompted). This step stops the appropriate services.

Examples

For examples of how to use dsinfo, refer to the examples in “[DSSTRT](#)” on page 402 (dsinfo, dsstop, and dsstrt have identical parameters and flags.)

DSSTRT

DSSTRT starts one or more services for one or more library systems on the server where it is run.

Usage

```
dsstrt -h
```

```
dsstrt -i library system -c components -g NO
```

Parameters and Switches

| | |
|----------------------|---|
| -h | Causes help text to be displayed. Cannot be used with other flags. |
| -i | Must precede the <i>librarysystem</i> parameter. |
| <i>librarysystem</i> | The name of the library system. If you do not specify a name, dsstrt starts the specified services for all library systems on the server where it is run. |
| -c | Must precede the <i>components</i> indicators. |
| <i>components</i> | You can include one or more of the following indicators to specify services you want to start. R Replication Service RA All Replication Services on the specified server S Storage Manager SA All Storage Managers on the specified server T Content Search Manager TA All Content Search Managers on the specified server |
| -g | When followed by NO, specifies no screen output. |

Usage Notes

- If you do not specify any component values, dsstrt starts all services except replication for the specified library system.
- Because flags are used to identify each parameter, you can type the parameters in any order, as long as the correct flags are included.

Procedure

To run dsstrt:

1. On the property or storage server, access the directory containing **dsstrt.exe**.

| OS Platform | Location |
|-------------|--|
| Windows | <i>StorageManagerInstallDrive:\filenet</i> <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Pbin</i> (for backward compatibility) |
| UNIX | <i>\$IDMDS_HOME</i> <i>StorageManagerInstallDir/filenet/librarysystem/pbin</i> (for backward compatibility) |

2. On the command line, enter:

```
dsstrt
```

and the desired parameters and flags. You must enter the parameters and flags on the command line (that is, you will not be prompted). This step starts the appropriate services, which may take several minutes to complete.

Examples

The following lists contain examples of dsstrt as it might be used to start library system services in a single or multi-system configuration. Using the same parameter syntax shown, you can use dsstop to stop library system services or dsinfo to get the status of those services.

Single-system Site

On machines that contain library system services for a single library system, you can use dsstrt to start system services, as shown in the following example list.

| Type this command: | To do this operation: |
|-----------------------|---|
| dsstrt | Start all system services (except Replication) on the machine where the utility is run. |
| dsstrt -i Acctng -c R | Start only the Replication service for the Acctng library system on the machine where the utility is run. (Note: You must specify the -i flag followed by the system name.) |
| dsstrt -c S | Start only the Storage Manager on the machine where the utility is run. |
| dsstrt -c T | Start only the Content Search Manager on the machine where the utility is run. |
| dsstrt -c ST | Start the Storage Manager and Content Search Manager on the machine where the utility is run. |

Multisystem Site

On machines that contain library system services for two or more library systems, you can use dsstrt to start system services, as shown in the following example list.

| Type this command: | To do this operation: |
|-----------------------|--|
| dsstrt | Start all services (except Replication) for all library systems on the machine where the utility is run. |
| dsstrt -i Acctng | Start all services (except Replication) for a system (Acctng) on the machine where the utility is run. |
| dsstrt -i Acctng -c R | Start the Replication service for a system (Acctng) on the machine where the utility is run. |
| dsstrt -i Acctng -c S | Start the Storage Manager for a system (Acctng) on the machine where the utility is run. |
| dsstrt -i Acctng -c T | Start the Content Search Manager for a system (Acctng) on the machine where the utility is run. |
| dsstrt -c RA | Start all Replication services on the machine where the utility is run. |
| dsstrt -c SA | Start all Storage Managers on the machine where the utility is run. |
| dsstrt -c TA | Start all Content Search Managers on the machine where the utility is run. |

DTCVLCHK

DTCVLCHK detects potential Document Class or CVL (Controlled Vocabulary List) collisions between two or more library systems in an administrative (i.e., replication) domain. See [“Administrative Domains” on page 606](#) for more information.

DTCVLCHK is a read-only utility. It does not modify any part of the database, and it is safe to run while the replication daemon is running. For the most accurate results, DTCVLCHK should not be run when any Document Class or CVL replication requests are pending. If it is run while such requests are pending, the only undesired result would be a message about the affected objects not existing or not matching.

You can use DTCVLCHK to detect problems in any of the following situations:

- Before they occur when bringing a new library system into an existing replication domain.
- To monitor a library system replication domain for Document Classes or CVLs being created on incorrect machines.

DTCVLCHK produces the following output:

Note Refer to [“Definition of Terms for USGRPCHK” on page 450](#) for the criteria USGRPCHK uses to determine whether objects are considered *effectively identical* for replication purposes.

- For each library in the domain, DTCVLCHK:
 - Creates a list of source and replica DocTypes.
 - Creates a list of source and replica CVLs.
 - Attaches a list of values to each CVL. Reports any values which do not have a corresponding CVL object (each CVL is represented by a unique CVL object in the library system).
- Checks all source DocTypes against each other, looking for duplicates by name. If the duplicate found is the DOCTYPE_CONTROL_OBJECT, reports whether they are not effectively identical. Otherwise, reports a duplicate DocType and indicates whether the duplicate source DocTypes are effectively identical or not.
- Checks all source CVLs against each other, looking for duplicates by name. Reports duplicates and indicates whether the duplicate source CVLs are effectively identical or not.
- For every replica DocType, extracts the source library system and checks for a matching source object on the indicated library system. Warns if the indicated library system doesn't exist, if the replica is not effectively identical to the source, or if there is no matching source DocType object.
- For every replica CVL:
 - Extracts the source library system and checks that there is a matching source on the indicated library system. Warns if the indicated library system does not exist, if the replica is not effectively identical to the master, or if there is no matching master.
 - Checks all the CVL values. Warns if the value is a source, or is a replica from a library system that is not the CVL's source library system.

- For every source DocType (that isn't DOCTYPE_CONTROL_OBJECT) and CVL, checks that every library system has a replica of the same name (except the library system that contains each master object). Displays a warning if no replica is found.
- Reports a list of DocTypes sourced on each library, for reference.

Usage

In all modes, DTCVLCHK requires Admin access to at least one library system. You can provide the administrative user name and password on the command line or be prompted for them. DTCVLCHK is a DOS command-line utility and can be invoked with any of the following options, depending on what information about the replication domain you want to see.

The following command lists all library systems in the administrative domain to which the specified library system belongs.

```
dtcvlchk -L domain_master_library domain_master_server [admin_user admin_pwd]
```

The following command performs collision and inconsistency detection among all reachable library systems in the administrative domain to which the specified library system belongs. It logs the output to *logfile*. The *-v* switch displays logged output on the screen as well. While running, the program will prompt for an admin username and password for the other library systems in the domain.

```
dtcvlchk -D [-v] domain_master_library domain_master_server logfile [admin_user admin_pwd]
```

The following command performs collision detection only between the two library systems. The *logfile* parameter and *-v* switch function are described above. This is the mode you would use when bringing a new library system into an existing administrative domain.

```
dtcvlchk [-v] domain_master_library domain_master_server new_library new_library_server logfile [admin_user1 admin_pwd1 admin_user2 admin_pwd2]
```

The following command lists doctypes and CVLs in the administrative domain that were not mastered on the specified library and server and reports the state of the replication trigger for each library in the domain.

```
dtcvlchk -P [-v|-q] domain_master_library domain_master_server logfile [admin_user admin_pwd]
```

Parameters and Switches

- | | |
|----|--|
| -D | Performs collision detection between all reachable library systems in the replication administrative domain to which <i>domain_master_library</i> belongs. If used with the <i>-v</i> switch, <i>-D</i> must precede <i>-v</i> . |
| -L | Lists all library systems in the replication administrative domain to which <i>domain_master_library</i> belongs. |
| -P | Lists doctypes and CVLs in the administrative domain that were not mastered on the specified library and server and reports the state of the replication trigger for each library in the domain. |

| | |
|--|---|
| <code>-v</code> | Displays logged output on the screen. Note Parameters are order-specific when using the [-v] option; the -D parameter must precede the -v. |
| <code>logfile</code> | A text file containing a report of the collisions, inconsistencies, and missing replicas among the various Document Class and CVL objects. |
| <code>domain_master_library</code> | The name of the domain master library system for the replication administrative domain. Refer to “ Replication Services Overview ” on page 594 for more information. |
| <code>domain_master_server</code> | The computer name of the server where the domain master library system's database is installed. |
| <code>new_library</code> | The name of the library system you are adding to the replication administrative domain. |
| <code>new_server</code> | The computer name of the server where the new library system's database is installed. |
| <code>admin_user</code> and <code>admin_pwd</code> | The user name and password for the library system Admin account. DTCVLCHK requires Admin access to the designated administration library system. You can specify the administrative user name and password on the command line; if you do not, DTCVLCHK will prompt you for them. Note The <code>admin_user</code> and <code>admin_pwd</code> parameters must be used as a pair; neither can be used on the command line without the other. Do not attempt to use these parameters on the command line if <code>admin_pwd</code> is null. The parameters <code>admin_user1</code> and <code>admin_pwd1</code> represent the user name and password for the first library system specified on the command line, and <code>admin_user2</code> and <code>admin_pwd2</code> represent the user name and password for the second library system specified on the command line. |

Definition of Terms for DTCVLCHK

Effectively Identical Document Classes

Document Classes are *effectively identical* if the following are the same:

- Name
- Property Details
 - Property
 - CVL Name
 - Reqd (Required: Yes/No)
 - Default Value

Note To compare these values, use the CS Document Class and CVL Admin, available via the CS Explorer Tools menu.

The property lists for Document Classes present special challenges when you try to compare them because they can contain standard properties (that have the same IDs in all library systems) and custom properties (that have potentially different IDs that must be mapped to one another across library systems). So, it is entirely possible to have the “same” custom property on multiple library systems with a different ID on each. For properties on two library systems to be identical:

- They must both reference the same property, whether it be a standard property with the same ID in both lists or a custom property that matches by ID mapping. System properties are always the same across library systems, but custom properties with the same name and mapped IDs are only considered *effectively identical* if they match in terms of:
 - Type
 - Length
 - Parent object (that is, whether it is an Item, Version, or Version Checkout)
 - Multivalue setting (Yes/No)

Note If you have custom properties of the same name that differ in the above attributes, to make them match you must rename one of them (both the custom property name and its language-specific display name), copy or create a new custom property that effectively matches the one on the other library system, and then modify the Document Class to reference the new custom property.

- They must both reference the same CVL, or not reference a CVL.
- They must both be required or not required.
- They must both have the same default value, or no default value.

Effectively Identical CVLs

CVLs are *effectively identical* if the following are the same:

- CVL Name
- CVL Data Type
- CVL Data

Note To compare these values, use the CS Document Class and CVL Admin, available via the CS Explorer Tools menu. Select a CVL Name and click Edit CVL to display the CVL Information. Refer to the *CS Admin Tools Online Help*.

DocType/CVL Object Collision

An object collision occurs when two DocType or CVL objects on different library systems have the same name and are *effectively identical* as defined above.

DocType/CVL Object Inconsistency

An object inconsistency occurs when two DocType or CVL objects on different library systems have the same name but are not *effectively identical* as defined above.

Procedure

You can run DTCVLCHK on a 32-bit Windows machine that has the appropriate database client software and Client Library Files installed and configured to connect to the necessary library systems associated with the replication administrative domain.

Note It will generally be most convenient to run DTCVLCHK on a server within the domain, since such a machine should have the correct Client Library Files and configuration.

To run DTCVLCHK:

1. Insert the *FileNetContent Services Admin Tools* CD.
2. Change to the *drive:\CLI53\util* directory on the CD.
3. From a DOS command prompt, execute DTCVLCHK using one of the command options described in the “Usage” topic.
4. Review the logfile and use the CS Admin Tools to reconcile any reported problems. Refer to the *CS Admin Tools Online Help* for procedures on changing property values.

DTCVLCHK Messages

Note In the following error messages:

- The term *option* is equivalent to *switch*.
- The term *master object* is equivalent to *source object*.
- The term *mastered* is equivalent to *sourced*.
- There are references to options (switches) that are intended for internal FileNet purposes only. Do not use any switches or parameters that are not listed in the “Parameters and Switches” topic.

-D mode

- SYSTEM OBJECT: DOCTYPE_CONTROL_OBJ is effectively identical between libraries *system1* on *server1* and *system2* on *server2*.
- SYSTEM OBJECT: DOCTYPE_CONTROL_OBJ is different between libraries *system1* on *server1* and *system2* on *server2*.

DOCTYPE_CONTROL_OBJECT holds certain pieces of information: the custom property ID used on items to hold their Document Classes, and the default DocType (Document Class) for new items. Even though the IDs of the custom properties may differ, that should be concealed by the custom property mapper. Generally, they should be identical except when different default DocTypes have been set for the two library systems, which is a legitimate and harmless difference.

- DUP MASTER: DocClass *name* is mastered on library *system1* on *server1* and on library *system2* on *server2*. The two masters are effectively identical.
- DUP MASTER: DocClass *name* is mastered on library *system1* on *server1* and on library *system2* on *server2*. The two masters differ.

This message also applies to CVLs. These messages will be common in upgraded library systems. Since neither DocTypes nor CVLs could be replicated in a version 4.3 library system, it was necessary to create them

on each system. As a consequence, when the domain is upgraded, there will be duplicate masters. On a newly installed 5.1 library system, or an upgraded library system after migration has been completed, this will usually mean that someone has created an object on the wrong library system. It could also mean that a migration failed during the movement of an object.

- HYBRID CVL: CVL *name* is mastered on library *system1* on *server1*, but value *value* within that CVL is a replica from library *system2*.
- HYBRID CVL: CVL *name* is mastered on library *system1*, but value *value* within that CVL is mastered on library *system2* on *server2*.

Hybrids occur when CVLs with the same name but differing values are created on two different library systems; the replication daemon merges the values.

- MISSING MASTER: The replica of DocClass *name* found on library *system1* on *server1* claims to be mastered on *system2*, but there is no such library in the domain.

This message also applies to CVLs. This is not likely to occur unless the database has been altered, or the database has become corrupted.

- WARNING: The replica of DocClass *name* found on library *system1* on *server1* claims to be mastered on *system2*, which we could not connect to. Cannot verify this DocClass.

This message also applies to CVLs. Occurs when the source library system for this DocType or CVL cannot be reached, so there is no way to check whether the replica is valid.

- CORRUPT REPLICA: The replica of DocClass *name* found on library *system1* on *server1* does not match its master on library *system2*.

This message also applies to CVLs. This could be due to replication failure or database corruption but it can likely be fixed by forcing the object to replicate again.

- MISSING MASTER: The replica of DocClass *name* found on library *system1* on *server1* claims to be mastered on *system2*, but there is no such DocClass on that library.

This message also applies to CVLs. This message indicates a serious problem. This message will also appear if *system2* has a replica of a DocClass or CVL object. Ideally, another master of the object can be located somewhere in the domain.

- MISSING REPLICA: There is no replica of DocClass *name* (mastered on library *system1* on *server1*) on library *system2* on *server2*.

This message also applies to CVLs. DTCVLCHK checks that for every source object, every other library system in the domain has a replica by the same name. It does not, however, check the replica's source library system because that has already been done elsewhere. This message will be seen frequently during upgrades because there are no existing replicas of any DocType or CVL.

Two-library Mode

- SYSTEM OBJECT: DOCTYPE_CONTROL_OBJ is effectively identical on the two libraries.
- SYSTEM OBJECT: DOCTYPE_CONTROL_OBJ is different on the two libraries.

DOCTYPE_CONTROL_OBJECT holds the following pieces of information: the custom property ID used on items to hold their Document Classes, and the default DocType (Document Class) for new items. Even though

the IDs of the custom properties may differ, that should be concealed by the custom property mapper. Generally, they should be identical except when different default DocTypes have been set for the two systems, which is a legitimate and harmless difference.

- SYSTEM OBJECT: DocTypes CVL is effectively identical on the two libraries.
- SYSTEM OBJECT: DocTypes CVL is different on the two libraries.

The CVL named “DocTypes” is a system object that lists the names of all DocTypes on that library system. Within an administrative domain, it should be identical on each library system. For library systems that are *not* in the same domain, it is okay for lists to be different.

- DUPLICATE: DocClass *name* is effectively identical on both libraries.
- DUPLICATE: DocClass *name* appears on both libraries but differs between them.

This message also applies to CVLs.

DTCVLMIG

DTCVLMIG migrates Document Classes and CVLs between two or more library systems. It is used to consolidate Document Class and CVL administration on one machine within a replication domain (which is the recommended configuration). The utility can be used to fix a library system domain that was configured incorrectly.

DTCVLMIG does minimal consistency checking. It assumes that you have run DTCVLCHK and cleaned up any problems. If the system is in an inconsistent state, running DTCVLMIG can greatly exacerbate the problem.

DTCVLMIG makes changes to the states of Document Classes and CVLs. The order in which the changes are made has been carefully chosen such that inadvertent interruption of the program, loss of network connectivity, and other errors should leave the system in a functional state with regard to Document Classes and CVLs. It may not be in a consistent state, but the system should be able to function while the administrator investigates and decides what actions to take. For example, running DTCVLCHK to determine the state of the library system.

If not specified, DTCVLMIG will prompt the user for the required admin user names and passwords. In all modes, DTCVLMIG requires Admin access rights to at least one library system.

For a given source object on library system A, there are seven possible cases that can occur on library system B (as described below). The number of each case corresponds to numbered options (or switches).

- 0 There is no object of the same name.
- 1 There is an identical source object of the same name.
- 2 There is a non-identical source object of the same name.
- 3 There is an identical replica object of the same name.
- 4 There is a non-identical replica object of the same name.
- 5 There is an identical replica object of the same name, but it has another library system (not library system A) as its source library system (that is, a “foreign” replica).
- 6 There is a non-identical replica object of the same name, but it has another library system (not library system A) as its source library system (that is, a “foreign” replica).

Note DTCVLMIG supports case 0 and 1. You cannot use DTCVLMIG to address cases 2 through 6.

Usage

Warning DTCVLMIG contains some switches that are intended for internal FileNet purposes only. Do not deviate from the following usage examples or use any switches or parameters that are not listed in the “Parameters and Switches” topic.

To list all library systems in the replication domain to which the specified library system belongs:

```
dtcvlmig -L domain_master_library domain_master_server [admin_user admin_pwd]
```

To add objects in the *current_home_library* to the *domain_master_library* and make the source object a replica of the new object in the *domain_master_library*:

```
dtcvlmig [-v] [-x] -0a current_home_library current_home_server  
domain_master_library domain_master_server logfile [admin_user1 admin_pwd1  
admin_user2 admin_pwd2]
```

To demote objects in the *current_home_library* to be replicas of objects in the *domain_master_library*:

```
dtcvlmig [-v] [-x] -1d current_home_library current_home_server  
domain_master_library domain_master_server logfile [admin_user1 admin_pwd1  
admin_user2 admin_pwd2]
```

To check and enable replication on all the libraries in the administrative domain:

```
dtcvlmig -R [-v|-q] domain_master_library domain_master_server logfile [admin_user  
admin_pwd]
```

To check and enable replication on just the library specified in the command line:

```
dtcvlmig -R1 [-v|-q] domain_master_library domain_master_server logfile [admin_user  
admin_pwd]
```

Parameters and Switches

| | |
|-------------------|--|
| -0a | Adds the object to the target library system and makes the source object a replica of the new target object. |
| -1d | Demotes the source object to be a replica of the target object. |
| -L | Lists all library systems in the replication domain. |
| -R | Checks and enables replication on all the libraries in the administrative domain. This is equivalent to setting "Replication Enabled" to "Yes" in the Admin Tools. |
| -R1 | Checks and enables replication on the specified library. This is equivalent to setting "Replication Enabled" to "Yes" in the Admin Tools. |
| -v | Displays logged output on the screen. |
| -x | Causes DTCVLMIG to migrate objects. If the -x switch is not specified, a listing of objects to be migrated is produced, but no migration is performed. |
| | Caution You may want to initially run this utility without the -x switch to check for anomalous messages and errors. |
| <i>admin_pwd</i> | Administrative user password for the specified library system. |
| <i>admin_pwd1</i> | Administrative user password for the first library system specified on the command line. |
| <i>admin_pwd2</i> | Administrative user password for the second library system specified on the command line. |

| | |
|------------------------------|--|
| <i>admin_user</i> | Administrative user name for the specified library system. |
| <i>admin_user1</i> | Administrative user name for the first library system specified on the command line. |
| <i>admin_user2</i> | Administrative user name for the second library system specified on the command line. |
| <i>current_home_library</i> | A library system other than the domain master where the Document Classes and/or CVLs are currently source objects instead of replicas. |
| <i>current_home_server</i> | A library system server other than the domain master server where the Document Classes and/or CVLs are currently source objects instead of replicas. |
| <i>domain_master_library</i> | The name of the domain master library system for the replication administrative domain. Refer to “Replication Services Overview” on page 594 for more information. |
| <i>domain_master_server</i> | The computer name of the server where the domain master library system's database is installed. |
| <i>logfile</i> | The name of a text file that contains the command's output. |

Procedure

You can run DTCVLMIG on a 32-bit Windows machine that has the appropriate database client software and Client Library Files installed and configured to connect to the necessary library systems associated with the replication administrative domain.

Note It will generally be most convenient to run DTCVLMIG on a server within the domain, since such a machine should have the correct Client Library Files and configuration.

To run DTCVLMIG:

1. Insert the *FileNet Content Services Admin Tools* CD.
2. Change to the *drive:\CLI53\util* directory on the CD.
3. From a DOS command prompt, execute DTCVLMIG using one of the command options specified in the “Usage” topic.
4. Review the logfile and use the CS Admin Tools to reconcile any reported problems. Refer to the *CS Admin Tools Online Help* for procedures.

DTCVLMIG Messages

You may see one or more of the following error messages when the library system is in two-library system mode:

Note In the following error messages:

- The term *homelib* is equivalent to *source library*.
- The term *option* is equivalent to *switch*.

- The term *master object* is equivalent to *source object*.
- There are references to options (switches) that are intended for internal FileNet purposes only. Do not use any switches or parameters that are not listed in the “Parameters and Switches” topic.
- *Library1* and *library2* must be different!
- *Library system* on *server* is not part of a domain. Both libraries must be part of a domain. Aborting.
- *Library source_system* on *source_server* is not in the same domain as *library target_system* on *target_server*. Both libraries must be in the same domain. Aborting.
- Moving master DocClasses from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -d switch is specified.

- Moving master DocClasses and CVLs from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -d or -c switch is specified.

- Moving master DocClasses and CVLs and CVL values from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -d, -c, -a, or no switch is specified.

- Moving master DocClasses and CVL values from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -d or -a switch is specified.

- Moving master CVLs from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -c switch is specified.

- Moving master CVLs and CVL values from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -c or -a switch is specified.

- Moving master CVL values from *library source_system* on *source_server* to *library target_system* on *target_server*.

Displayed when the -a switch is specified.

- Running in TEST mode. No objects will be moved.

The -x switch was not specified.

- Running in EXECUTE mode! Objects will be moved!

The -x switch was specified.

- DocClasses have been modified.

CVLs have been modified.

CVL values have been modified.

At least one object has been successfully or partially migrated.

- Replication triggers of *system* on *server* successfully disabled.
- Migrating DocClasses from *source_system* on *source_server* to *target_system* on *target_server*.

A similar message is displayed for CVLs and CVL values.

- Replication triggers of *system* on *server* successfully re-enabled.
- WARNING: DocType *name* does not exist on the target library. It will be added.

Displayed when no options are specified (the default). A similar message is displayed for CVLs and CVL values.

- WARNING: DocType *name* does not exist on the target library. No action will be taken for this DocType.

Displayed when the *-On* option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: An identical master of DocType *name* exists on the target library. The source will be made into a replica of the existing target.

Displayed when the *-ld* option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: An identical master of DocType *name* exists on the target library. The source will overwrite the target.

Displayed when the *-lo* option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: An identical master of DocType *name* exists on the target library. No action will be taken for this DocType.

Displayed when *-ln* or the default option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing master of DocType *name* exists on the target library. The source will be made into a replica of the existing target.

Displayed when the *-2d* option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing master of DocType *name* exists on the target library. The source will overwrite the target.

Displayed when the *-2o* option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing master of DocType *name* exists on the target library. No action will be taken for this DocType.

Displayed when the default option is specified. A similar message is displayed for CVLs and CVL values.

- An identical replica of DocType *name* exists on the target library. Normal migration.

Displayed when the default option is specified. A similar message is displayed for CVLs and CVL values.

- An identical replica of DocType *name* exists on the target library. No action will be taken for this DocType.

Displayed when the `-3n` option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing replica of DocType *name* exists on the target library. The source will overwrite the target.

Displayed when the `-4o` option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing replica of DocType *name* exists on the target library. No action will be taken for this user.

Displayed when the default option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: An identical replica of DocType *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will be made into a replica with the same homelib as the target.

Displayed when the `-5d` option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: An identical replica of DocType *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will overwrite the target.

Displayed when the `-5o` or `-h` switch is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: An identical replica of DocType *name* exists on the target library, but the replica's homelib (*system*) is not the source library. No action will be taken for this DocType.

Displayed when the default option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing replica of DocType *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will be made into a replica with the same homelib as the target.

Displayed when the `-6d` option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing replica of DocType *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will overwrite the target.

Displayed when the `-6o` option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: A differing replica of DocType *name* exists on the target library, but the replica's homelib (*system*) is not the source library. No action will be taken for this user.

Displayed when the default option is specified. A similar message is displayed for CVLs and CVL values.

- WARNING: Bad internal type for DocType *name*! No action will be taken for this DocType.

Internal DTCVLMIG program error.

- Some custom properties are missing on the target library. Cannot add this DocType.

DTCVLMIG scans the source DocType's property list looking for custom properties and attempts to map them to custom properties on the target library system, in the same way that the replication daemon does when replicating DocTypes. This message means that it could not map all of the custom properties found in the source's property list, so this DocType (which does not exist on the target) will not be added there.

- Some custom properties are missing on the target library. Cannot move this DocType.

Same description as the previous message, except the DocType did exist on the target and will be overwritten.

- DocType's added-by user *source_user* does not exist on target, changing to *target_admin_user*.

DocType's owner *source_user* does not exist on target, changing to *target_admin_user*.

These messages mean the same thing. Similar messages are displayed for CVLs and CVL values.

- Cannot move matching CVLDATA for DocType *name* because it does not exist on source library.

DTCVLMIG couldn't add the CVLDATA to the target library system because it couldn't find it on the source library system.

When a DocType does not exist on the target library system, DTCVLMIG adds the DocType and the CVLDATA in CVL "DocTypes." If DTCVLMIG did not do this, the CS Document Class and CVL Admin would add them, which would result in multiple masters or other undesirable situations.

- Error: Cannot migrate CVL value *value* in *CVLname* to target library because CVL *CVLname* does not exist on target!

Trying to migrate a CVL value, but there is no appropriate CVL object on the target.

- Error: Cannot migrate CVL value *value* in *CVLname* to target library because CVL *CVLname* on target library has differing datatype.

Trying to migrate a CVL value, but the CVL object on the target has a different datatype.

FILE2ID

Using the specified encrypted file name for a version, this utility generates the associated item ID and version ID for the version file.

Usage

```
file2id encrypted_filename
```

Parameters and Switches

encrypted_filename The specified encrypted file for a version, as identified in the storage repository.

Usage Notes

Use ID2FILE to do the opposite operation of FILE2ID, that is, to generate the version's encrypted file name from that version's item ID and version ID.

Procedure

To run FILE2ID:

1. On the initial storage server, access the directory containing the **file2id.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

2. On the command line, enter:

```
file2id encrypted_filename
```

This example displays the item ID and version ID of the version on your monitor.

HMSRVCHK

The HMSRVCHK utility checks the replication folder on which you want to enable Home Server. It generates a list of all subfolders and documents that will be sourced on Home Server on running the HMSRVMIG utility and the documents and subfolders that have exceptions. The subfolders and documents with exceptions will not be sourced on Home Server.

Note It is advisable to run the HMSRVCHK utility before running the HMSRVMIG utility. This will enable you to know about all the documents and subfolders that will be considered exceptions. Subsequently, you can resolve these exceptions before running the HMSRVMIG utility to ensure that the required documents and subfolders are sourced on Home Server.

Refer to the “HMSRVMIG” on page 424 to know more about the HMSRVMIG utility.

Exceptions

The HMSRVCHK utility considers the following documents and subfolders as exceptions:

- The documents and subfolders that are sourced on an outside domain library.
- The documents that have been filed in a replication folder, which is not Home Server enabled.
- Any document participating in a compound document as parent or child.
- A document or a subfolder with pending work request(s) for which the Replication Daemon has tried maximum number of times to complete the request(s) and failed.

Usage

You can run HMSRVCHK with any of two modes, depending on what information you want it to produce. In all modes, HMSRVCHK requires Admin access to home server library system. You can provide the administrative user name and password on the command line, or be prompted for them.

The following command lists all target library systems for the specified replication folder.

```
hmsrvchk -L domain_master_library domain_master_server root_folder_ID [admin_user  
admin_pwd]
```

The following command performs discrepancy detection for all the documents and subfolders within the specified replication folder. It logs the output to *logfile*.

```
hmsrvchk -D [-q] domain_master_library domain_master_server root_folder_ID logfile  
[admin_user admin_pwd]
```

Parameters and Switches

- | | |
|----|--|
| -L | Lists all domain member libraries and target libraries to which the specified folder is replicating. |
| -D | Performs discrepancy detection among all the target library systems in the administrative domain for the specified folder. |

| | |
|---------------------------------------|--|
| <code>-q</code> | Turn off the screen logging. Note Parameters are order-specific when using the [-q] switch; the -D parameter must precede the -q. |
| <code>domain_master_library</code> | The name of the Domain Master library system for the replication administrative domain. See the section “ Replication Services Overview ” on page 594 for more information. |
| <code>domain_master_server</code> | The computer name of the server where the domain master library system's database is installed. |
| <code>root_folder_id</code> | ID of the replication folder on the Domain Master on which you want to enable Home Server. |
| <code>logfile</code> | A text file containing a report of the discrepancies among the various documents and subfolders within the specified replication folder. |
| <code>admin_user and admin_pwd</code> | The User name and password for the library system Admin account. HMSRVCHK requires Admin access to the Domain Master library system. You can specify the administrative user name and password on the command line; if you do not, HMSRVCHK will prompt you for them. Note The <code>admin_user</code> and <code>admin_pwd</code> parameters must be used as a pair; neither can be used on the command line without the other. Do not attempt to use these parameters on the command line if <code>admin_pwd</code> is null. |

Procedure

You can run the HMSRVCHK utility on a 32-bit Windows machine that has the appropriate database client software and Client Library Files installed and configured to connect to the necessary library systems associated with the replication administrative domain.

Note It will generally be most convenient to run HMSRVCHK on a server within the domain, since such a machine should have the correct Client Library Files and configuration.

To prepare to run HMSRVCHK:

- Ensure that all member libraries participating in replication of the specified replication folder are available and connected to the Domain Master library.
- To ensure that you get accurate results, run the utility when there are no pending work requests for any document or subfolder within the replication folder.

To run HMSRVCHK:

1. Insert the *FileNet Content Services Admin Tools* CD.
2. Change to the `drive:\CLI54\util` directory on the CD.

3. From a DOS command prompt, execute HMSRVCHK using one of the command options described in the "Usage" topic.
4. Review the logfile and resolve exceptions for subfolders and documents which you want to source to Home Server.

HMSRVCHK Messages

You can see one or more of the following messages while running the HMSRVCHK utility:

- Username or password is incorrect for *domain master library* on *domain master server*. Please try again.
Displayed when the login fails. User would be prompted to enter user credentials again.
- Only administrative group members can run this utility. Please try again.
Displayed when Login profile belongs to non-administrative group.
- The given library *Domain Master* is not a domain master. Cannot apply home server on this library.
Displayed when the specified library system on which you are running the utility is not the Domain Master.
- The given folder ID *root_folder_ID* is incorrect for *domain master library* on *domain master server*. Please try again.
Displayed when the specified folder ID is not of a root folder on the Domain Master.
- Home server cannot be applied on the root folder *root_folder_ID* since it is replicating only to outside domain libraries.
Displayed when the specified replication folder is not replicating to any member library.
- Error *errorCode* connecting to *library name* on *server name*.
Displayed when one of the target libraries is not available/connected to the Domain Master.

Discrepancy log file

On running the HMSRVCHK utility, a discrepancy log file is generated. This log file contains a list of all the subfolders and documents within the replication folder which will be:

- considered exception and will not be sourced on Home Server.
- selected to be sourced on Home Server by swapping the status of the object between the member library and Home Server.

The log file also provides the reasons for the exceptions found in the replication folder.

If the log file with the specified name is already present, then the existing log file is overwritten.

Log file messages

The log file contains the following messages:

- `<date> <time> -- [FOLDER EXCEPTION] Folder "folder name" [folder ID] is source on Outside Domain library [Outside Domain library name] with ID [ID]. It will not be sourced to domain master.`

Displayed when the folder is sourced on an outside domain library and will not be sourced.

- `<date> <time> -- Folder "folder name" [folder ID] is source on library [library name] with ID [ID]. It will be sourced to domain master.`

Displayed when the folder will be successfully sourced on Home Server.

- `<date> <time> -- [ITEM EXCEPTION] Item "Item name" [Item ID] is source on library [library name] with ID [ID]. It is also filed in non home server enabled folder. It will not be sourced to domain master.`

Displayed when the item is filed in a non Home Server enabled folder. As a result, the item will not be sourced on Home Server.

- `<date> <time> -- [ITEM EXCEPTION] Item "Item name" [Item ID] is source on library [library name] with id [ID]. It is participating in compound document relationship. It will not be sourced to domain master.`

Displayed when the item is a part of compound document as parent or child. As a result, the item will not be sourced on Home Server.

- `<date> <time> -- [ITEM EXCEPTION] Item "Item name" [Item ID] is source on library [library name] with id [ID]. It has pending replication request with maximum error count reached. It will not be sourced to domain master.`

Displayed when the item has any pending work requests. As a result, the item will not be sourced on Home Server.

- `<date> <time> -- Item "Item name" [Item ID] is source on library [library name] with ID [ID]. It will be sourced to domain master.`

Displayed when the item will be successfully sourced on Home Server.

HMSRVMIG

The HMSRVMIG utility enables you to source all the documents and subfolders within a replication folder from member libraries to Home Server. This utility is useful when the source of the documents and subfolders within a replication folder is scattered across member libraries and the Domain Master. Using this utility, you can source all these documents and subfolders to Home Server to ensure centralized management.

Note Refer to “Replication Home Server” on page 616 and “Enabling Replication Home Server” on page 625 to know more about the Replication Home Server feature for which this utility is used.

Usage

You can run the HMSRVMIG utility in any of the following two modes depending upon the desired outcome:

To list all target libraries for the specified replication folder

```
hmsrvmig -L domain_master_library domain_master_server root_folder_ID [admin_user  
admin_pwd]
```

To source subfolders and documents within the specified replication folder from all target member libraries on the specified Home Server

```
hmsrvmig -D [-q] domain_master_library domain_master_server root_folder_ID logfile  
[admin_user admin_pwd]
```

Parameters and switches

| | |
|------------------------------|---|
| -L | Lists all domain member libraries and target libraries to which the specified folder is replicating. |
| -D | Sources subfolders and documents from all target libraries in the administrative domain to the specified Home Server. |
| -q | Turn off the screen logging. Note Parameters are order-specific when using the [-q] switch; the -D parameter must precede the -q. |
| <i>domain_master_library</i> | The name of the domain master library system for the replication administrative domain. See the section “Replication Services Overview” on page 594 for more information. |
| <i>domain_master_server</i> | The computer name of the server where the domain master library system's database is installed. |
| <i>root_folder_id</i> | ID of the replication folder on the Domain Master on which you want to enable Home Server. |
| <i>logfile</i> | The name of a text file that contains the command's output. |

admin_user and admin_pwd

The User name and password for the library system Admin account.

HMSRVMIG requires Admin access to the Domain Master library system. You can specify the administrative user name and password on the command line; if you do not, HMSRVMIG will prompt you for them.

Note The *admin_user* and *admin_pwd* parameters must be used as a pair; neither can be used on the command line without the other. Do not attempt to use these parameters on the command line if *admin_pwd* is null.

Procedure

You can run the HMSRVMIG utility on a 32-bit Windows machine that has the appropriate database client software and Client Library Files installed and configured to connect to the necessary library systems associated with the replication administrative domain.

Note It will generally be most convenient to run HMSRVMIG on a server within the domain, since such a machine should have the correct Client Library Files and configuration.

To prepare to run HMSRVMIG:

- Ensure that all member libraries participating in replication of the specified replication folder are available and connected to the Domain Master library.
- Make sure the "Logons Enabled?" property is set to "No" for all member libraries participating in the replication of the specified replication folder.
- Use the ReplMon utility on all member libraries to ensure that there are no pending work requests for any document or folder that has to be migrated.
- (Recommended) Run the Home Server Check (HMSRVCHK) utility to check all the documents and subfolders within the replication folder for exceptions.
- Resolve the exceptions for the documents and subfolders that you want to source to Home Server.
- Make backups of all the member libraries participating in replication.

Note Refer to the section "HMSRVCHK" to know more about the exceptions.

To run the HMSRVMIG utility

1. Insert the *FileNet Content Services Admin Tools* CD.
2. Change to the *drive:\CLI54\util* directory on the CD.
3. From a DOS command prompt, execute HMSRVMIG using one of the command options specified in the "Usage" topic.
4. Review the logfile.

HMSRVMIG Messages

You can see one or more of the following messages while running the HMSRVMIG utility:

- Username or password is incorrect for *domain master library* on *domain master server*. Please try again.
Displayed when the login fails. User would be prompted to enter user credentials again.
- Only administrative group members can run this utility. Please try again.
Displayed when Login profile belongs to non-administrative group.
- The given library *Domain Master* is not a domain master. Cannot apply home server on this library.
Displayed when the specified library system on which you are running the utility is not the Domain Master.
- The given folder ID *root_folder_ID* is incorrect for *domain master library* on *domain master server*. Please try again.
Displayed when the specified folder ID is not of a root folder on the Domain Master.
- Home server cannot be applied on the root folder *root_folder_ID* since it is replicating only to outside domain libraries.
Displayed when the specified replication folder is not replicating to any member library.
- Error *errorCode* connecting to *library name* on *server name*.
Displayed when one of the target libraries is not available/connected to the Domain Master.
- Logon not disabled at library *domain member library* on *domain member server*. Please disable the logon and run utility again.
Displayed when user logons are enabled at one or more target libraries.
- There are some pending item / folder replication request at library *domain member library* on *domain member server*. Run replication services to ensure that all the replication entries are processed before running the utility.
Displayed when there are pending work request(s) existing on one or more target libraries.

Migration log file

On running the HMSRVMIG utility, a migration log file is generated. This log file contains a list of all the subfolders and documents within the replication folder which have been:

- considered exception and have not been sourced on Home Server.
- successfully sourced to Home Server by swapping the status of the object between the member library and Home Server.

If the log file with the specified name is already present, then the existing log file is overwritten.

Log file messages

The log file contains messages related to:

- `<date> <time> -- [FOLDER EXCEPTION] Folder "folder name" [folder ID] is source on Outside Domain library [Outside Domain library name] with ID [ID]. It is not sourced to domain master.`

Displayed when the folder has not been sourced on Home Server because the folder is sourced on an outside domain library.

- `<date> <time> -- Folder "folder name" [folder ID] is source on library [library name] with ID [ID]. It is sourced to domain master.`

Displayed when the folder has been successfully sourced on Home Server.

- `<date> <time> -- [ITEM EXCEPTION] Item "Item name" [Item ID] is source on library [library name] with ID [ID]. It is also filed in non home server enabled folder. It is not sourced to domain master.`

Displayed when the item is filed in a non Home Server enabled folder. As a result, the item has not been sourced on Home Server.

- `<date> <time> -- [ITEM EXCEPTION] Item "Item name" [Item ID] is source on library [library name] with id [ID]. It is participating in compound document relationship. It is not sourced to domain master.`

Displayed when the item is a part of compound document as parent or child. As a result, the item has not been sourced on Home Server.

- `<date> <time> -- [ITEM EXCEPTION] Item "Item name" [Item ID] is source on library [library name] with id [ID]. It has pending replication request with maximum error count reached. It is not sourced to domain master.`

Displayed when the item has any pending work requests. As a result, the item has not been sourced on Home Server.

- `<date> <time> -- Item "Item name" [Item ID] is source on library [library name] with ID [ID]. It is sourced to domain master.`

Displayed when the item has been successfully sourced on Home Server.

ID2FILE

ID2FILE uses the specified item ID and version ID for a given version file to generate the encrypted file name for that version.

Usage

```
id2file item_id version_id
```

Parameters and Switches

| | |
|-------------------|--|
| <i>item_id</i> | The item ID used to generate the encrypted file name. |
| <i>version_id</i> | The version ID used to generate the encrypted file name. |

Usage Notes

Use FILE2ID to do the opposite operation of ID2FILE, that is, to generate the version's item ID and version ID from that version's encrypted file name.

Procedure

To run ID2FILE:

1. On the initial storage server, access the directory containing the **id2file.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

2. On the command line, enter:

```
id2file item_id version_id
```

This example displays the encrypted file name of the version on your monitor.

INDXDATE

Description

INDXDATE indexes for content search those versions that were last accessed (that is, checked out, checked in, or copied) before a specified date.

Usage

```
indxdate /username= /password= /propertyserver=  
/systemname= /date= /indextype= /itemsearch=  
/listonly= /logfile= /inputfile=
```

Parameters and Switches

| | |
|------------------|--|
| /username= | The user name of the person indexing the versions. Only versions to which this user has at least Owner access rights can be indexed. To avoid potential access violations, this utility should in most cases be run by the Admin user. |
| /password= | The password, if any, for the user identified with the <i>username</i> parameter. If you use a batch file, we suggest that you not include the <i>password</i> parameter in your script (it can be read by anyone who can view the contents of that file). Instead, you should allow the utility to prompt you for this value (your response will be masked on the screen as you type it). |
| /propertyserver= | The name of the server where the database engine is installed. |
| /systemname= | The name of the library system. |
| /date= | The identifying version date. Versions that were last accessed before 12:00 a.m. of this date will be flagged for indexing. This date must appear in the format <i>mm/dd/yyyy</i> (for example, 06/11/1995). |
| /indextype= | The type of indexing operation to perform: I Index R Reindex D Deindex |
| /itemsearch= | The name of a pre-existing item search. Note that INDXDATE will produce inaccurate results if any search type other than an item search is specified for this parameter. Only versions of items returned by this search may be indexed (based on the <i>date</i> parameter). If you do not include this optional <i>itemsearch</i> parameter, INDXDATE will simply look at <i>all</i> documents last accessed before the <i>date</i> parameter to determine what to index. |
| /listonly= | A Yes/No value to indicate whether you want to run the utility in list mode only (the default value is Yes). |

`/logfile=` The name of the file to which you want the log of the INDXDATE results sent. This log will identify the following information of the versions that were indexed: Item ID, Version ID, Last access date, and number of bytes freed per version.

You can specify a value for this parameter in one of three ways:

- Accept the default value, **indxdate.log**. (If this file already exists, the new information will be appended to it.)
- Specify a different log file name. (If the file already exists, the new information will be appended to it.)
- Use `con` as the parameter value to indicate that you want the log printed to your monitor rather than to a file (for example, `logfile=con`).

`/inputfile=` The name of the file containing prompt-response values for the command line parameters.

Usage Notes

- On the command line, you must include a forward slash (/) before each parameter and a blank space between parameter statements. No space should appear between the parameter and its value; however, the parameter values themselves may contain blank spaces.

For example, you might type (on one line):

```
indxdate /username=John Braughn  
/propertyserver=01 /systemname=MZ1 /date=06/11/1997  
/itemsearch=mysrch /listonly=n
```

and press Enter to start the utility. If you leave any parameters out (as in this example), you will be prompted for them individually.

- You must run INDXDATE from the initial storage server or an additional storage server.
- The first property from your search is the default sort order setting. (An exception to this is the keyword property which cannot be used to sort search results.) If item ID is not the first property in your stored search, you must change the sort order.
- (MS-SQL only) On average, indexing 1000 versions will consume 70 MB of transaction log space. We recommend setting your log full threshold to 80% or less and indexing versions in increments so that the log does not fill to capacity. After you run INDXDATE, we suggest you perform a full system backup to save the content search repositories and to clear the log.

Procedure

To run INDXDATE from the command line:

1. If you want to use a search to identify the versions to index, you must create that search (and save it at the System access level) before running INDXDATE. Use CS Explorer or any other interface to the library system.
2. On the initial storage server or additional storage server, access the directory containing the **indxdate.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

3. On the command line, enter:

```
indxdate
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

The search, if any, is run first. Then, as the index operations take place, the results are either displayed on your monitor or printed to the specified log file.

To run **INDXDATE** from a batch file:

1. If you want to use a search to identify the versions to index, you must create that search (and save it at the System access level) before running **INDXDATE**. Use CS Explorer or any other interface to the library system.
2. Create a batch file with the following line in it:

```
indxdate /inputfile=indxdate.res
```

We recommend against including your library system password in the batch file. (Instead, allow the system to prompt you for it.) However, if you prefer to include the password so that this operation can be run completely unattended, replace the preceding batch file command with:

```
indxdate /password=value /inputfile=indxdate.res
```

If your library system does not have a password, type in instead:

```
indxdate /password= /inputfile=indxdate.res
```

3. Create the **indxdate.res** file. In this file, type in your values for the command line parameters, excluding password and inputfile, in the correct order as shown in the following example. Include a space before and after equal signs (=) and separate each command line parameter with a carriage return. For example:

```
username = John Doe
propertyserver = server01
systemname = system01
date = 06/11/1995
indextype = I
itemsearch = indxdate search
listonly = no
logfile = indxdate.log
```

4. When you are ready to run **INDXDATE**, start the batch file.

The search, if any, is run first. Then, as the index operations take place, the results are either displayed on your monitor or printed to the specified log file.

INDXFIX

INDXFIX searches for and queues for indexing all versions with an indexing status of “filter failed” or “index failed”. The search is created by INDXFIX as a search for all versions in the database.

INDXFIX should be run whenever you install new indexing filters or whenever problems with a Storage Manager are suspected of causing indexing failures.

Note The utility REPCHECK does not repair index failures or filter failures.

Usage

```
indxfix /username= /password= /propertyserver=  
/systemname= /currentversiononly= /listonly= /logfile=  
/inputfile=
```

Parameters and Switches

| | |
|----------------------|--|
| /username= | The user name of the person indexing the versions. Only versions to which this user has at least Owner access rights can be indexed. To avoid potential access violations, this utility should in most cases be run by the Admin user. |
| /password= | The password, if any, for the user identified with the <i>username</i> parameter. If you use a batch file, we suggest that you not include the <i>password</i> parameter in your script (it can be read by anyone who can view the contents of that file). Instead, you should allow the utility to prompt you for this value (your response will be masked on the screen as you type it). |
| /propertyserver= | The name of the server where the database engine is installed. |
| /systemname= | The name of the library system. |
| /currentversiononly= | A Yes/No value to indicate whether you want to index only the current version or all versions up to the indexed version limit. |
| /listonly= | A Yes/No value to indicate whether you want to run the utility in list mode only (the default value is Yes). |
| /logfile= | The name of the file to which you want the log of the INDXFIX results sent. This log will identify the following information of the versions that were indexed: Item ID, Version ID, Last access date, and number of bytes freed per version. You can specify a value for this parameter in one of three ways: <ul style="list-style-type: none">• Accept the default value, indxfix.log. (If this file already exists, the new information will be appended to it.)• Specify a different log file name. (If the file already exists, the new information will be appended to it.)• Use con as the parameter value to indicate that you want the log printed to your monitor rather than to a file (for example, logfile=con). |

`/inputfile=` Optional: The name of the file containing prompt-response values for the command line parameters.

Usage Notes

- On the command line, you must include a forward slash (/) before each parameter and a blank space between parameter statements. No space should appear between the parameter and its value, however, the parameter values themselves may contain blank spaces.

For example, you might type (on one line):

```
indxfix /username=John Braughn /propertyserver=01 /systemname=MZ1 /  
currentversiononly=y /listonly=n
```

and press Enter to start the utility. If you leave any parameters out (as in this example), you will be prompted for them individually.

- You must run INDXFIX from the initial storage server or an additional storage server.
- On average, indexing 1000 versions will consume 70 MB of transaction log space. We recommend setting your log full threshold to 80% or less and indexing versions in increments so that the log does not fill to capacity. After you run INDXFIX, we suggest you perform a full system backup to save the content search repositories and to clear the log.

Procedure

To run INDXFIX from the command line:

- On the initial storage server or additional storage server, access the directory containing the **indxfix.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

- On the command line, enter:

```
indxfix
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

INDXFIX runs a version search to find all the versions with indexing status of “filter failure” or “index failure”. Then, for each version found to be indexed, the results are either displayed on your monitor or printed to the specified log file.

To run INDXFIX from a batch file:

- Review the INDXFIX *Usage Notes* topic.
- Create a batch file with the following line in it:

```
indxfix /inputfile=indxfix.res
```

We recommend against including your library system password in the batch file (rather, allow the system to prompt you for it). However, if you would prefer to include the password so that this operation can be run completely unattended, you must replace the preceding batch file command with the following line:

```
indxfix /password=value /inputfile=indxfix.res
```

To indicate that you have no password, type in instead:

```
indxfix /password= /inputfile=indxfix.res
```

3. Create the **indxfix.res** file. In this file, type in your values for the command line parameters, excluding password and inputfile, in the correct order as shown in the following example. Separate each command line parameter with a carriage return.

```
username = John Doe  
propertyserver = server01  
systemname = system01  
currentversiononly = yes  
listonly = no  
logfile = indxfix.log
```

4. When you are ready to run INDXFIX, start the batch file.

As the index operations take place, the results are either displayed on your monitor or printed to the specified log file.

LDAP_EXP

Binds to an LDAP directory, searches for requested group names, writes an XML file of user/group information to be parsed by the CS_IMPORT utility, and closes the LDAP directory. Before running LDAP_EXP, you must install the shared files that it needs (see “Installation” on page 436). Refer to “Importing Users and Groups” on page 272 for information on using CS_IMPORT and LDAP_EXP.

Usage

```
ldap_exp /s server /p port /d bindDN /w bindPW -t ldap-server-type
  /b baseDN /g group-list-filename /o XML-output-filename /l logfile
  /x group-list-filename -c1 user_object_class -c2 group_object_class
  -a1 user_attribute -a2 group_attribute -a3 member_attribute /v /? /h
```

Parameters and Switches

| | |
|----------------------------|--|
| <i>server</i> | The name of the LDAP server or IP address where the LDAP directory is installed. |
| <i>port</i> | (Optional) LDAP port for the server where the LDAP directory is installed. The default is 389. |
| <i>bindDN</i> | LDAP Administrator bind distinguished name for binding to the LDAP directory. |
| <i>bindPW</i> | Encrypted LDAP Administrator bind password (see Usage Notes entry below for directions on password encryption) |
| <i>ldap-server-type</i> | The type of LDAP server: <ul style="list-style-type: none"> • <i>msft</i> Microsoft Active Directory • <i>nov</i> Novell eDirectory (NDS) • <i>sun</i> Sun ONE Directory (iPlanet) • <i>ibm</i> IBM Directory Server |
| <i>baseDN</i> | (Optional) LDAP base distinguished name. The part of the LDAP directory held on the local server. |
| <i>group-list-filename</i> | (Optional) Pathname of input file containing line-feed terminated group names. |
| <i>XML-output-filename</i> | (Optional) Pathname of the output XML file. The file must <i>not</i> be on a mapped drive. |
| <i>logfile</i> | (Optional) Pathname of logfile. |
| <i>group-list-filename</i> | (Optional) Pathname of output file containing group names |
| <i>user object class</i> | Class type for User. Default is ‘person’ for every ldap-server-type |
| <i>group object class</i> | Class type for Group. Default is ‘groupofuniquenames’ for sun ‘groupofnames’ for nov ‘group’ for msft |

| | |
|-------------------------|--|
| <i>user attribute</i> | Attribute to use for user. Default is 'userprincipalname' for msft. 'uid' for sun. 'cn' for nov. |
| <i>group attribute</i> | Attribute to use for group. Default is 'cn' for every ldap-server-type |
| <i>member attribute</i> | Attribute for member. Default is 'member' for msft and nov. 'uniquemember' for sun. |
| v | (Optional) Generate verbose log output. |
| ? | Generates simple program-execution help text. |
| h | Generates simple program-execution help text. |

Installation

Before running LDAP_EXP, you need to install the LDAP shared library files on the machine where you will be invoking the utility.

To install the files, copy the shared library files (those files whose names are of the form *.dll) from the *CS Admin Tools* CD-ROM, at <CD-ROM_Drive>:\util\ldap\NT32I, to a directory in the search path stored in the Path system environment variable—for example, the directory where the Content Services Client Library Files are installed: <CS_Shared_Files_Drive>:\Program Files\FileNet\Shared

Usage Notes

- LDAP_EXP is strictly a command-line oriented program; and all its required parameters are passed to it via the command line.
- You can use services of the operating system to automatically run LDAP_EXP at specified intervals. In Windows 2000 or Windows 2003, use the Schedule service. In Unix, use the cron command.
- Each parameter and switch is required unless otherwise noted.
- The command switches may be prefixed by either '-' or '/'.
- To obtain bindPW, your encrypted password, invoke the command-line utility pwcrypt, located in the filenet\systemname\pbin directory, as follows:

```
pwcrypt plaintext_password
```

The output of the pwcrypt command is

```
"plaintext_password" password is "encrypted_password" encrypted.
```

Example:

```
pwcrypt foobar
```

```
"foobar" password is "<,=8//" encrypted.
```

- LDAP_EXP first tries to use the "Group of Unique Names" attribute (which is itself a Distinguished Name attribute) to distinguish user/group names. The program then tries to use the "Group of Names" attribute to

distinguish user/group names and eliminates duplicates. Finally, the program tries to organize user/group names based on the Organizational Unit attribute.

- LDAP_EXP skips user names in Active Directory that have no Universal Principal Name. (The Universal Principal Name is the user's logon name combined with his domain. For example, a user named *user1* in domain *domain1* has UPN *user1@domain1.com*.)

Procedure

To run LDAP_EXP:

1. Review the LDAP_EXP *Usage Notes* topic.
2. Access the directory containing LDAP_EXP.

| OS Platform | Location |
|-------------|--|
| Windows | <i>CSMInstallDrive:\filenet\LibrarySystem\util</i> |
| UNIX | <i>\$IDMDS_HOME/librarysystem/util</i> |

3. On the command line, invoke the command, as in the following example:

```
ldap_exp -s SunSrvr1 -d "cn=Directory Manager" -w ":/, .++=" -t sun -b "dc=iPlanet,  
dc=com"
```

NLSADMIN

NLSADMIN installs a display language for a library system.

Usage

```
nlsadmin -h
```

```
nlsadmin -s <server> -i <catalog> -a -f <filename> -v
```

Parameters and Switches

| | |
|------------------------|--|
| <i>property_server</i> | The name of the server where the library services are installed. |
| <i>catalog</i> | The name of the library system to which you are adding a language. |
| <i>-a</i> | Add a language. |
| <i>filename</i> | The pathname to the NLS file for the display language to be added. |
| <i>-v</i> | Generate verbose output. |
| <i>-h</i> | Generate simple program-execution help text. Cannot be used with other switches. |

Usage Notes

- You can choose from the following display languages that ship with Content Services: Chinese, English, French, German, Japanese, and Korean (**chinese.nls**, **english.nls**, **french.nls**, **german.nls**, **japanese.nls**, and **korean.nls**). Other language files can be obtained with direct FileNet assistance.

Procedure

To run NLSADMIN:

- Verify that the Oracle processes are running.
- Log on to a storage server of the library system to which you intend to add a display language.
- Access the directory containing **nlsadmin**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

- On the command line, enter:

```
nlsadmin
```

and the desired parameters and switches.

Example

To install a French display language on a server named *server1* for library system named *library1*:

```
nlsadmin -s server1 -i library1 -a -f ..\nls\french.nls
```

REPCHECK

REPCHECK verifies that the logical links between the property repository (database) and the storage or content search repositories are intact. REPCHECK can either report inconsistencies or attempt to fix them.

Caution For Microsoft SQL Server databases, running REPCHECK in fix mode has the potential to fill your transaction log. When you run REPCHECK in fix mode, be sure to use a well-defined limit on the max number of errors to fix. Also, you may need to adjust your transaction log configuration in SQL Server. One way to avoid this problem is to enable the “Truncate log on checkpoint” feature before running REPCHECK (and then disable it afterward if desired). This option is available in Enterprise Manager, on the Options tab of the database properties sheet. Another way to avoid the problem is to configure the transaction log file to auto-grow. This option is available in Enterprise Manager, on the Transaction Log tab of the database properties sheet. If you use this method, you should monitor the process to ensure that the transaction log file does not fill up your disk. Refer to *SQL Server Books Online* for more information about these database features.

Usage

```
repcheck property_server librarysystem repository_type  
repository_id max_errors check_mode log_path
```

Parameters and Switches

| | |
|------------------------|---|
| <i>property_server</i> | The name of the server where the Property Manager is installed. |
| <i>librarysystem</i> | The name of the library system. |
| <i>repository_type</i> | The kind of repository. The options are: c Content Search Manager s Storage Manager |
| <i>repository_id</i> | The ID number of the storage repository or the content search repository (as identified in the Storage Repository or Content Search Repository object). |
| <i>max_errors</i> | The maximum number of errors that will be fixed or reported before the tool will stop and exit. |
| <i>check_mode</i> | The repcheck operation mode: r Report only f Fix c Create New Index and Fix (content search repository only) |
| <i>log_path</i> | The directory path where the log file will be written. This parameter is optional. The default location is the directory where the utility is located. |

Usage Notes

- We recommend that you periodically run REPCHECK and check for files in the **lost+found** directories.

For the library system to work properly, the integrity of the links between the storage and content search repositories, and the Property Manager (database) must be strictly maintained. Over time, however,

discrepancies may arise. For example, users may receive error messages indicating that the library system cannot find a version file or that a file has been found that does not have a corresponding object in the property repository. Such conditions may be caused by power outages during particular operations or by internal software problems.

- Because fixing a large number of inconsistencies may take a long time, we recommend running REPCHECK in the Report mode first. You can then use the *max_errors* parameter to adjust the number of inconsistencies that will be fixed each time REPCHECK is run.
- When running REPCHECK in the Fix mode against a storage repository, REPCHECK will not attempt to fix items that have been in the storage repository less than 10 minutes. However, it will report any inconsistencies found in the new items. This feature prevents REPCHECK from attempting to fix items that are in the process of being added to a library system.
- REPCHECK must be run from the initial storage server or an additional storage server, and the library system services must be running. On Windows platforms, REPCHECK must have access to the 32-bit client DLLs.
- REPCHECK can be run in the following modes:
 - Report only: Creates a report.
 - Fix: Creates a report and attempts to fix any inconsistencies.
- When the repository type is Storage Manager, REPCHECK will create a report named **stmgrxxx.log** (where **xxx** is the repository ID) in the directory containing the **repcheck** executable or in the directory specified by the *log_path* parameter. The log file (for example, **stmgr001.log**) will contain the run date and time, the check mode, and a table identifying any inconsistencies. The table will contain the following columns:
 - Condition: A code identifying the type of inconsistency. See below for code descriptions.
 - Item/Version: The item ID and version ID of the affected item and version.
 - Filename: The name of the affected shelf file (as identified in the **shelfxxx** directory).
 - User: The name of the user who checked in the item version.
 - Fix Status: The status code returned by REPCHECK. For example, a return status of 0 indicates success and a return status of 16477 indicates that a general DBI error has been detected. Refer to the “[Status Codes](#)” on page 573 for a description of available status codes.
- When the repository type is Content Search Manager, REPCHECK will create a report named **csmgrxxx.log** (where **xxx** is the repository ID) in the directory containing the **repcheck** executable or in the directory specified by the *log_path* parameter. The log file (for example, **csmgr001.log**) will contain the run date and time, the check mode, and a table identifying any inconsistencies. The table will contain the following columns:
 - Condition: A code identifying the type of inconsistency. See below for code descriptions.
 - Item/Version: The item ID and version ID of the affected item and version.
 - Index ID: The content index ID (CID) of the item version.

- **Fix Status:** The status code returned by REPCHECK. For example, a return status of 0 indicates success and a return status of 16477 indicates that a general DBI error has been detected. Refer to the “[Status Codes](#)” on page 573 for a description of available status codes.
- The following table describes the condition codes and the action that will be taken if REPCHECK is run in either Fix or Create New Index and Fix mode.

| Code | Description | Fix Mode Action |
|------|---|--|
| 1 | Invalid file on the storage repository. | Move file to directory lost+found located one directory up from storage repository directory. |
| 2 | Version in property repository indicates file is located in the storage repository, but it isn't located there. | Mark the version status as Online Lost. |
| 3 | Version in property repository points to file on different storage repository and version file exists in this repository. | Update version storage repository ID to value of storage repository in check. |
| 4 | File exists in the storage repository with no corresponding version in the property repository. | Move file to directory lost+found located one directory up from storage repository directory. |
| 5 | Version file status is Offline when file exists in the storage repository. | Mark version file status as Online. |
| 6 | Version file status is marked as Offline - NonReclaimable when file exists in the storage repository. | Mark version file status as Online. |
| 7 | Version file status is marked as Online Lost when file exists in the storage repository. | Mark version file status as Online. |
| 8 | Version file status is marked as Offline Lost when file exists in the storage repository. | Mark version file status as Online. |
| 9 | Version entry is indexed in content search repository but version doesn't exist in property repository. | Remove entry from content search repository. |
| 10 | Version in property repository exists but isn't marked as indexed. The version does exist in the content search repository. | Remove entry from content search repository. |
| 11 | Version in property repository is marked as being indexed but it isn't in the content search repository. | Request is made to reindex this version. |

| Code | Description | Fix Mode Action |
|------|--|---|
| 12 | Version in property repository is marked as being indexed and entry exists in content search repository, but the catalog entry contains an incorrect pointer reference to the version. | Entry is removed from content search repository and the version is reindexed. |
| 13 | Version in the property repository is marked as being indexed but points to an incorrect reference to the content search repository, meanwhile the content search repository shows this version has been properly indexed. | Entry is removed from content search repository and the version is reindexed. |

- A library system administrator must take the appropriate action to address each file that appears in the **lost+found** directory. For example, the administrator may need to add the file to the appropriate storage repository (making the necessary modifications to the Property Manager) or she may need to delete the file.

Procedure

To run REPCHECK:

1. On the initial storage server, access the directory containing the **repcheck.exe**.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

2. On the command line, enter:

```
repcheck
```

and, optionally, one or more of the REPCHECK parameters in the order specified above. You will be prompted for the remaining parameters that you did not include on the command line (except for the optional *log_path* parameter). To specify an alternative path for the REPCHECK log file, you must enter all of the parameters (including *log_path*) on the command line.

3. Respond to the REPCHECK results. Refer to the REPCHECK *Usage Notes* topic for details.

REPLKICK

REPLKICK forces Users, Groups, DocClasses, and CVLs to replicate throughout a replication administrative domain (see “Administrative Domains” on page 598 for a detailed description)

REPLKICK is a command-line utility for Windows. You run it against the domain master library system to synchronize library systems. You can run it on any machine that has the appropriate Client Library Files installed and configured to connect to the necessary property servers. Typically, it is most convenient to run it on the administrative domain master server.

Usage

```
replkick [-u] [-g] [-d] [-c] [-v] [-f] domain_master_library domain_master_server  
logfile [admin_user admin_pwd]
```

Parameters and Switches

| | |
|---|--|
| <code>-u</code> | Explicitly replicate users. |
| <code>-g</code> | Explicitly replicate groups. |
| <code>-d</code> | Explicitly replicate DocClasses. |
| <code>-c</code> | Explicitly replicate CVLs. |
| <code>-v</code> | Echo the logfile output to the screen. |
| <code>-f</code> | Force REPLKICK to run on non-domain master. |
| <code><i>domain_master_library</i></code> | The library system name of the administrative domain master. If the library you specify is not the administrative domain master, an error is reported and REPLKICK quits. |
| <code><i>domain_master_server</i></code> | The computer name where the domain master library system is installed. |
| <code><i>logfile</i></code> | File name for the log generated by this utility. |
| <code><i>admin_user</i></code> | The username for the library system Admin account. REPLKICK requires admin access to the designated administration library. You can specify the administrative username and password on the command line; if you do not, REPLKICK will prompt you for them. |
| <code><i>admin_pwd</i></code> | The password for the library system Admin account. |

Usage Notes

All four object types (users, groups/memberships, DocClasses, and CVLs) are queued for replication.

You can safely run REPLKICK while repldaemon is active. Even if repldaemon scans the table while REPLKICK is running, that just means that some raw requests get converted now and some will have to wait until next time repldaemon scans the table.

Repldaemon uses the same work ID for all three of DOCTYPE/CVLOBJ/CVLDATA, so they all show up in replmon as the DocClass requests; this is correct behavior.

Whenever repldaemon receives a DOCTYPE ACL request, it automatically transfers all ACLs belonging to the given DOCTYPE, so only one request is needed, not one per ACL.

Procedure

You must run REPLKICK on a 32-bit Windows machine that has the appropriate Client Library Files installed and configured to connect to the necessary property servers associated with the replication administrative domain. It is usually most convenient to run REPLKICK on the administrative domain master server or from a library system administrative client (for example, in an HP-UX server environment).

To run REPLKICK:

1. Insert the *FileNet Content Services Admin Tools* CD.
2. Change to the *drive:\CLI53\util* directory on the CD.
3. On the command line, enter:

```
replkick
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line.

Example

The following is an annotated example of REPLKICK run against two library systems: BUNNIES and SKUNKS. For this example, REPLKICK was run without any object type options specified (-u, -g, -d, or -c), so all object types were queued.

```
Queueing USERS for replication:
```

```
  Chandra  
  Syndi
```

```
Queueing GROUPS for replication:
```

```
  Hares
```

Note: *This group has no members.*

```
  Lapines
```

```
    member Chandra  
    member Syndi
```

```
Queueing DocTypes for replication:
```

```
  Experiment  
  General
```

```
Queueing CVLOBJs for replication:
```

```
  DocTypes
```

```
Queueing CVLDATAs for replication:
```

```
  DocTypes Experiment  
  DocTypes General
```

Note: *There is actually a third value in the DocTypes CVL, but it is a replica from SKUNKS so it was not queued.*

SPISTAT

SPISTAT prints the programming defined value for a specific error code.

Note For a listing and description of all status codes refer to “Status Codes” on page 573.

Usage

```
spistat status_code
```

Parameters and Switches

status_code The status code value.

Procedure

To run SPISTAT:

1. On the initial storage server, access the directory containing the **spistat** executable.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

2. On the command line, enter:

```
spistat status_code
```

This displays the define value for the status code on your monitor.

SRDELETE

SRDELETE deletes items found by a given search.

Usage

```
srdelete propertyserver librarysystem username password search /delete:limit
```

Parameters and Switches

| | |
|-----------------------|---|
| <i>propertyserver</i> | The name of the server where the Property Manager is installed. |
| <i>librarysystem</i> | The name of the library system. |
| <i>username</i> | The user name for logging in to the specified library system server. |
| <i>password</i> | The password for the specified user. |
| <i>search</i> | The name of the saved item search to perform. Note that SRDELETE will produce inaccurate results if any search type other than an item search is specified for this parameter. Multi-word search names must be enclosed in quotation marks. |
| <i>/delete:</i> | Indicates that the documents found by the search operation are to be deleted. |
| <i>limit</i> | The limit on the number of items to delete. This number can only be specified in conjunction with the <i>/DELETE</i> option, and must be preceded by a <i>:</i> (colon). |

Usage Notes

- You can specify any stored search request for SRDELETE, but it must be saved at the System access level, not at the User or Group access level. This requirement is intended to limit the use of SRDELETE to members of the Administrators group, which is the only group that can save searches at the System access level.
- The stored search request specified by SRDELETE must have item ID in its sort order. Note that the first property from your search is the default sort order setting. (An exception to this is the keyword property which cannot be used to sort search results.) If item ID is not in the sort order of your stored search, you must change the sort order. Refer to *CS Explorer Online Help* for details.
- If any parameters are omitted, you will be prompted for them. You can run SRDELETE from the command line or in batch mode (if all the arguments are specified). Note that this reveals your password in the batch file or script.
- SRDELETE deletes the items it finds only if you specify the */delete* parameter; otherwise SRDELETE simply lists the items that it would delete, without actually deleting them.
- You can force SRDELETE to stop after a certain number of items by including a *limit* parameter value after the */delete* parameter. For example, to stop after attempting to delete 2000 items, you would include the option: */delete:2000*.
- SRDELETE reports its activity to your screen and to two log files (in a tab delimited format). The **srdelete.log** file shows the items deleted, their number of versions, total disk space freed, and a result code. The **srfailed.log** file is similar, but is a log of items that failed to delete. This **srfailed.log** file lets you address unusual cases easily.

- SRDELETE renames old log files with a sequential extension, as in **srdelete.001**, **srdelete.002**, and so on. As many as 10 old log files can be preserved in this way. SRDELETE preserves both the **srdelete.log** and **srfailed.log** files.

Procedure

To run SRDELETE:

1. Review the SRDELETE *Comments* topic and *Examples* topic.
2. On the initial storage server, access the directory containing the **srdelete** executable.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

3. On the command line, enter:

```
srdelete
```

and, optionally, one or more of the associated parameters. You will be prompted for any parameters that you did not include on the command line. SRDELETE deletes the items it finds only if you specify the `/delete` parameter; otherwise SRDELETE simply lists the items that it would delete, without actually deleting them.

Examples

```
srdelete srvr01 mycat admin secretpwd olditems
```

```
srdelete srvr01 mycat admin "" "Unaccessed 120 days" /delete
```

```
srdelete srvr01 mycat admin secretpwd olditems /delete:1000
```

USGRPCHK

USGRPCHK detects potential user or group collisions between two or more library systems in a replication administrative domain (see [“Administrative Domains”](#) on page 606 for a detailed description).

USGRPCHK is a command-line utility for Windows. You can run it on library system administrative client or any machine that has the appropriate Client Library Files installed and configured to connect to the necessary library systems. USGRPCHK is a read-only utility; it does not modify any part of the database. It is safe to run while the replication daemon is running. For the most accurate results, do not run USGRPCHK when any user, group, or group membership replication requests are pending. If it is run while such requests are pending, the only undesired result would be a message about the affected users or groups not existing or not matching.

You can use USGRPCHK to detect problems in any of the following situations:

- Before they occur, when bringing a new library system into an existing administrative domain.
- To monitor a library system administrative domain for users or groups being created on incorrect machines.

USGRPCHK produces the following output:

Note Refer to [“Definition of Terms for USGRPCHK”](#) on page 450 for the criteria USGRPCHK uses to determine whether objects are considered *effectively identical* for replication purposes.

- Creates a list of source and replica users and groups.
- Attaches a list of group members to each group. Reports any memberships for non-existent groups or users and any groups that have no members.
- Checks all source users against each other, looking for duplicates by name. If the duplicate found is the Admin user or ReplDaemon, warns if they are not effectively identical. Otherwise, reports a duplicate user and whether the users are effectively identical.
- Checks all source groups against each other, looking for duplicates by name. If the duplicate found is the Administrators or General Users group, warns if they are not effectively identical. Otherwise, reports a duplicate group and whether the groups are effectively identical.
- For every replica user and group, extracts the source library system and checks that there is a matching source object on the indicated library system. Warns if the indicated library system doesn't exist, if the replica is not effectively identical to the source, or if there is no matching source object.
- For every source user and group, checks that every library system has a replica of the same name (except the library system that the source resides on). Displays a warning if no replica is found.
- Reports a list of users and groups sourced on each library system for reference.

Usage

You can run USGRPCHK with any of four modes, depending on what information you want it to produce. In all modes, USGRPCHK requires Admin access to at least one library system. You can provide the administrative user name and password on the command line, or be prompted for them.

The following command lists all library systems in the administrative domain to which the specified library system belongs.


```
usgrpchk -L domain_master_library domain_master_server [admin_user admin_pwd]
```

The following command performs collision and inconsistency detection among all reachable library systems in the administrative domain to which the specified library system belongs. It logs the output to *logfile*. The *-v* switch displays logged output on the screen as well. While running, the program will prompt for an admin user name and password for the other library systems in the domain.

```
usgrpchk -D [-v] domain_master_library domain_master_server logfile [admin_user admin_pwd]
```

The following command performs collision detection only between the two library systems. The *logfile* parameter and *-v* switch function are described above. This is the mode you would use when bringing a new library system into an existing administrative domain.

```
usgrpchk [-v] domain_master_library domain_master_server new_library new_library_server logfile [admin_user1 admin_pwd1 admin_user2 admin_pwd2]
```

The following command lists users and groups in the administrative domain that were not mastered on the specified library and server and reports the state of the replication trigger for each library in the domain.

```
usgrpchk -P [-v|-q] domain_master_library domain_master_server logfile [admin_user admin_pwd]
```

Parameters and Switches

| | |
|------------------------------|--|
| -D | Performs collision detection between all reachable library systems in the replication administrative domain to which <i>domain_master_library</i> belongs. If used with the <i>-v</i> switch, <i>-D</i> must precede <i>-v</i> . |
| -L | Lists all library systems in the replication administrative domain to which <i>domain_master_library</i> belongs. |
| -P | Lists users and groups in the administrative domain that were not mastered on the specified library and server and reports the state of the replication trigger for each library in the domain. |
| -v | Displays logged output on the screen. Note Parameters are order-specific when using the <i>[-v]</i> switch; the <i>-D</i> parameter must precede the <i>-v</i> . |
| <i>logfile</i> | A text file containing a report of the collisions, inconsistencies, and missing replicas among the various User and Group objects. |
| <i>domain_master_library</i> | The name of the domain master library system for the replication administrative domain. See “ Replication Services Overview ” on page 594 for more information. |
| <i>domain_master_server</i> | The computer name of the server where the domain master library system's database is installed. |
| <i>new_library</i> | The name of the library system you are adding to the replication administrative domain. |

new_library_server

The computer name of the server where the new library system is installed.

admin_user and *admin_pwd*

The user name and password for the library system Admin account.

USGRPCHK requires Admin access to the designated administration library system. You can specify the administrative user name and password on the command line; if you do not, USGRPCHK will prompt you for them.

Note The *admin_user* and *admin_pwd* parameters must be used as a pair; neither can be used on the command line without the other. Do not attempt to use these parameters on the command line if *admin_pwd* is null.

The parameters *admin_user1* and *admin_pwd1* represent the user name and password for the first library system specified on the command line, and *admin_user2* and *admin_pwd2* represent the user name and password for the second library system specified on the command line.

Definition of Terms for USGRPCHK

Effectively Identical User Objects

Two User objects are *effectively identical* if they have the same values for the following properties (each '?' that appears in the following list is part of the property name):

- User Name
- Default Group
- Default Path
- Default Storage Category
- Default Archive Category
- Add Items?
- Network Logon Name
- Primary Storage Manager
- Enabled?
- Default Index Version?
- Access Domain
- Display Language
- Secure Searches?
- User Full Name

- Content Search Threshold
- Imported User Type

Two Group objects are *effectively identical* if they have the same values for the following properties (each '?' that appears in the following list is part of the property name):

- Group Name
- Default Storage Category
- Default Archive Category
- Primary Storage Manager
- Enabled?
- Default Index Versions?
- Access Domain
- Display Language
- Secure Searches?
- Content Search Threshold
- Imported Group Type
- Group Memberships (the lists for the two groups being compared must be identical)

User and Group Properties Not Compared

The following properties are not compared when determining whether User or Group objects are *effectively identical*:

- Date Added (varies between source and replica)
- Comment (varies between source and replica; replicas always have a comment of the form [source library])
- Added By User (varies between source and replica; replicas are always owned by ReplDaemon)

User/Group Object Collision

An object collision occurs when two User or Group objects on different library systems have the same name and are *effectively identical* as defined above.

User/Group Object Inconsistency

An object inconsistency occurs when two User or Group objects on different library systems have the same name but are not *effectively identical* as defined above.

Procedure

You can run USGRPCHK on a 32-bit Windows machine that has the appropriate database client software and Client Library Files installed and configured to connect to the necessary library systems associated with the replication administrative domain.

Note It will generally be most convenient to run USGRPCHK on a server within the domain, since such a machine should have the correct Client Library Files and configuration.

To run USGRPCHK:

1. Insert the *FileNet Content Services Admin Tools* CD.
2. Change to the *drive:\CLI53\util* directory on the CD.
3. From a DOS command prompt, execute USGRPCHK using one of the command options described in the “Usage” topic.
4. Review the logfile and use the Content Services Admin Tools to reconcile any reported problems. Refer to the *Content Services Admin Tools Online Help* for procedures on changing property values.

Note USGRPCHK reports the fact that two users are not “effectively identical,” but it does not identify the User properties whose values differ.

USGRPCHK Messages

Note In the following error messages:

- The term *master* is equivalent to *source*.
- The term *mastered* is equivalent to *sourced*.
- The term *home library* is equivalent to *source library*.

-D mode

- WARNING: system user *name* differs between library *system1* on *server1* and library *system2* on *server2*.

WARNING: system group *name* differs between library *system1* on *server1* and library *system2* on *server2*.

While ReplDaemon is an internal user and should theoretically have the same name on each library system, the Admin user can legitimately have a different name on each library system, as can the Administrators and the General Users groups.

- DUP MASTER: User *name* is mastered on library *system1* on *server1* and on library *system2* on *server2*. The two masters are effectively identical.

DUP MASTER: Group *name* is mastered on library *system1* on *server1* and on library *system2* on *server2*. The two masters are effectively identical.

DUP MASTER: User *name* is mastered on library *system1* on *server1* and on library *system2* on *server2*. The two masters differ.

DUP MASTER: Group *name* is mastered on library *system1* on *server1* and on library *system2* on *server2*. The two masters differ.

These messages usually mean that a user or group has been created with the same name on two different library systems. This is not recommended, but it is not a serious condition. The consequence is that the other library systems in the domain may have a replica from one or the other, possibly depending on which one was changed most recently. (See a definition of *effectively identical* elsewhere in this topic.)

- MISSING MASTER: The replica of user *name* found on library *system1* on *server1* claims to be mastered on *system2*, but there is no such library in the domain.

MISSING MASTER: The replica of group *name* found on library *system1* on *server1* claims to be mastered on *system2*, but there is no such library in the domain.

These messages may occur if the database has been modified directly or if the database has become corrupted.

- WARNING: The replica of user *name* found on library *system1* on *server1* claims to be mastered on *system2*, which we could not connect to. Cannot verify this user.

WARNING: The replica of group *name* found on library *system1* on *server1* claims to be mastered on *system2*, which we could not connect to. Cannot verify this group.

These messages may occur when the source library system for this user or group cannot be reached. Consequently, it will not be possible to check whether the replica is valid.

- CORRUPT REPLICA: The replica of user *name* found on library *system1* on *server1* does not match its master on library *system2*.

CORRUPT REPLICA: The replica of group *name* found on library *system1* on *server1* does not match its master on library *system2*.

These messages may indicate a Replication failure or a corrupted database. Although serious, this condition can probably be rectified by forcing the object to replicate again.

- MISSING MASTER: The replica of user *name* found on library *system1* on *server1* claims to be mastered on *system2*, but there is no such user on that library.

MISSING MASTER: The replica of group *name* found on library *system1* on *server1* claims to be mastered on *system2*, but there is no such group on that library.

These messages indicate a serious problem. Ideally, a usable master of the object can be located somewhere in the domain. These messages will also appear if *system2* has a replica of object *name*. This can occur after migrating objects before using REPLKICK to propagate the new home library system information throughout the domain.

- MISSING REPLICA: There is no replica of user *name* (mastered on library *system1* on *server1*) on library *system2* on *server2*.

MISSING REPLICA: There is no replica of group *name* (mastered on library *system1* on *server1*) on library *system2* on *server2*.

During the domain check, USGRPCHK checks to see that for every source object, every other library system has a replica of the same name (although it does not check the replica's source library system; that has already been checked elsewhere). These messages indicate that no replica was found on library *system2*.

Two-library Mode

- DUPLICATE: User *name* is effectively identical on both libraries.

DUPLICATE: Group *name* is effectively identical on both libraries.

These messages indicate that effectively identical users or groups were added to the library systems, possibly before the administrative domain was created. (Refer to “[Definition of Terms for USGRPCHK](#)” on page 450 for a definition of effectively identical users and groups.)

- DUPLICATE: User *name* appears on both libraries but differs between them.

DUPLICATE: Group *name* appears on both libraries but differs between them.

These messages indicate that users or groups with the same name but with different values for at least one property were added to the library systems. This may have been done before the administrative domain was created. (Refer to “[Definition of Terms for USGRPCHK](#)” on page 450 for a definition of effectively identical users and groups.)

-D or Two-library Mode

- WARNING: In library *system1* on *server1*, the group name has no associated memberships.

A group with no members is unlikely (not recommended) but permitted. This message may be evidence of an incomplete action or an object that has not been completely deleted. Further investigation is warranted.

- WARNING: In library *system1* on *server1*, user *username* is a member of group *groupname*, but there is no such group.

WARNING: In library *system1* on *server1*, group *groupname* has nonexistent user *username* as a member.

These messages indicate a database integrity violation. The database may be seriously corrupted.

USGRPMIG

USGRPMIG migrates users and groups between two or more library systems. It is used to consolidate user and group administration on one machine within a replication domain (which is the recommended configuration). The utility can be used to fix a library system domain that was configured incorrectly.

USGRPMIG does minimal consistency checking. It assumes that you have run USGRPCHK and cleaned up any problems. If the library system is in an inconsistent state, running USGRPMIG can greatly exacerbate the problem.

USGRPMIG makes changes to the states of users and groups. The order in which the changes are made has been carefully chosen such that inadvertent interruption of the program, loss of network connectivity, and other errors should leave the library system in a functional state with regard to users and groups. It may not be in a consistent state, but the library system should be able to function while the administrator investigates and decides what actions to take. For example, running USGRPCHK to determine the state of the library system, and looking at the affected users and groups with the Content Services Admin Tools.

Certain users (the Admin user, ReplDaemon) and groups (the Administrators and General Users groups) are never replicated or migrated.

If not specified, USGRPMIG will prompt the user for the required admin user names and passwords. In all modes, USGRPMIG requires Admin access rights to at least one library system.

For a given source object on library system A, there are seven possible cases that can occur on library system B (as described below). The number of each case corresponds to numbered options (or switches).

- 0 There is no object of the same name.
- 1 There is an identical source object of the same name.
- 2 There is a non-identical source object of the same name.
- 3 There is an identical replica object of the same name.
- 4 There is a non-identical replica object of the same name.
- 5 There is an identical replica object of the same name, but it has another library system (not library system A) as its source library system (that is, a “foreign” replica).
- 6 There is a non-identical replica object of the same name, but it has another library system (not library system A) as its source library system (that is, a “foreign” replica).

Note USGRPMIG supports case 0 and 1. You cannot use USGRPMIG to address cases 2 through 6.

Usage

Warning USGRPMIG contains some switches that are intended for internal FileNet purposes only. Do not deviate from the following usage examples or use any switches or parameters that are not listed in the “Parameters and Switches” topic.

To list all library systems in the replication domain to which the specified library system belongs:

```
usgrpmig -L domain_master_library domain_master_server [admin_user admin_pwd]
```

To add objects in the *current_home_library* to the *domain_master_library* and make the source object a replica of the new object in the *domain_master_library*:

```
usgrpmig [-v] [-x] -0a current_home_library current_home_server  
domain_master_library domain_master_server logfile [admin_user1 admin_pwd1  
admin_user2 admin_pwd2]
```

To demote objects in the *current_home_library* to be replicas of objects in the *domain_master_library*:

```
usgrpmig [-v] [-x] -1d current_home_library current_home_server  
domain_master_library domain_master_server logfile [admin_user1 admin_pwd1  
admin_user2 admin_pwd2]
```

To check and enable replication on all the libraries in the administrative domain:

```
usgrpmig -R [-v|-q] domain_master_library domain_master_server logfile [admin_user  
admin_pwd]
```

To check and enable replication on just the library specified in the command line:

```
uspgrpmig -R1 [-v|-q] domain_master_library domain_master_server logfile [admin_user  
admin_pwd]
```

Parameters and Switches

| | |
|-------------------|---|
| -0a | Adds the object to the target library system and makes the source object a replica of the new target object. |
| -1d | Demotes the source object to be a replica of the target object. |
| -L | Lists all library systems in the replication domain. |
| -R | Checks and enables replication on all the libraries in the administrative domain. This is equivalent to setting "Replication Enabled" to "Yes" in the Admin Tools. |
| -R1 | Checks and enables replication on the specified library. This is equivalent to setting "Replication Enabled" to "Yes" in the Admin Tools. |
| -v | Displays logged output on the screen. |
| -x | Causes USGRPMIG to migrate objects. If the -x switch is not specified, a listing of objects to be migrated is produced, but no migration is performed. Caution You may want to initially run this utility without the -x switch to check for anomalous messages and errors. |
| <i>admin_pwd</i> | Administrative user password for the specified library system. |
| <i>admin_pwd1</i> | Administrative user password for the first library system specified on the command line. |
| <i>admin_pwd2</i> | Administrative user password for the second library system specified on the command line. |

| | |
|------------------------------|---|
| <i>admin_user</i> | Administrative user name for the specified library system. |
| <i>admin_user1</i> | Administrative user name for the first library system specified on the command line. |
| <i>admin_user2</i> | Administrative user name for the second library system specified on the command line. |
| <i>current_home_library</i> | A library system other than the domain master where the users and/or groups are currently source objects instead of replicas. |
| <i>current_home_server</i> | A library system server other than the domain master server where the users and/or groups are currently source objects instead of replicas. |
| <i>domain_master_library</i> | The name of the domain master library system for the replication administrative domain. See “ Replication Services Overview ” on page 594 for more information. |
| <i>domain_master_server</i> | The computer name of the server where the domain master library system's database is installed. |
| <i>logfile</i> | The name of a text file that contains the command's output. |

Procedure

You can run USGRPMIG on a 32-bit Windows machine that has the appropriate database client software and Client Library Files installed and configured to connect to the necessary library systems associated with the replication administrative domain.

Note It will generally be most convenient to run USGRPMIG on a server within the domain, since such a machine should have the correct Client Library Files and configuration.

1. Insert the *FileNet Content Services Admin Tools* CD.
2. Change to the *drive:\CLI53\util* directory on the CD.
3. From a DOS command prompt, execute USGRPMIG using one of the command options specified in the “Usage” topic.
4. Review the logfile and use the Content Services Admin Tools to reconcile any reported problems. Refer to the *Content Services Admin Tools Online Help* for procedures.

USGRPMIG Messages

You may see one or more of the following error messages when the library system is in two-library system mode:

Note In the following error messages:

- The term *homelib* is equivalent to *source library*.
- The term *option* is equivalent to *switch*.
- The term *master object* is equivalent to *source object*.

- There are references to options (switches) that are intended for internal FileNet purposes only. Do not use any switches or parameters that are not listed in the “Parameters and Switches” topic.
- *Library1* and *library2* must be different!
- Library *system* on *server* is not part of a domain. Both libraries must be part of a domain. Aborting.
- Library *source_system* on *source_server* is not in the same domain as library *target_system* on *target_server*. Both libraries must be in the same domain. Aborting.
- Moving master users and groups from library *source_system* on *source_server* to library *target_system* on *target_server*.

Displayed when the default options are specified.

- Running in TEST mode. No objects will be moved.

The -x switch was not specified.

- Running in EXECUTE mode! Objects will be moved!

The -x switch was specified.

- Replication triggers of *system* on *server* successfully disabled.
- Migrating user from *source_system* on *source_server* to *target_system* on *target_server*.

A similar message is displayed for groups.

- Users have been modified.

Groups have been modified.

At least one object has been successfully or partially migrated.

- Replication triggers of *system* on *server* successfully re-enabled.
- WARNING: user *name* does not exist on the target library. It will be added.

Displayed when no options are specified. A similar message is displayed for groups.

- WARNING: user *name* does not exist on the target library. No action will be taken for this User.

Displayed when the -On option is specified. A similar message is displayed for groups.

- WARNING: An identical master of user *name* exists on the target library. The source will be made into a replica of the existing target.

Displayed when the -ld option is specified. A similar message is displayed for groups.

- WARNING: An identical master of user *name* exists on the target library. No action will be taken for this user.

Displayed when no options are specified.

- WARNING: A differing master of user *name* exists on the target library. The source will be made into a replica of the existing target.
Displayed when the -2d option is specified. A similar message is displayed for groups.
- WARNING: A differing master of user *name* exists on the target library. The source will overwrite the target.
Displayed when the -2o option is specified. A similar message is displayed for groups.
- WARNING: A differing master of user *name* exists on the target library. No action will be taken for this user.
Displayed when no options are specified. A similar message is displayed for groups.
- An identical replica of user *name* exists on the target library. Normal migration.
Displayed when -3o, -3, or default options are specified. A similar message is displayed for groups.
- An identical replica of user *name* exists on the target library. No action will be taken for this user.
Displayed when the -3n option is specified. A similar message is displayed for groups.
- WARNING: A differing replica of user *name* exists on the target library. The source will overwrite the target.
Displayed when the -4o option is specified. A similar message is displayed for groups.
- WARNING: A differing replica of user *name* exists on the target library. No action will be taken for this user.
Displayed when no options are specified. A similar message is displayed for groups.
- WARNING: An identical replica of user *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will be made into a replica with the same homelib as the target.
Displayed when the -5d option is specified. A similar message is displayed for groups.
- WARNING: An identical replica of user *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will overwrite the target.
Displayed when the -5o option is specified. A similar message is displayed for groups.
- WARNING: An identical replica of user *name* exists on the target library, but the replica's homelib (*system*) is not the source library. No action will be taken for this user.
Displayed when the default, -5n or no options are specified. A similar message is displayed for groups.
- WARNING: A differing replica of user *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will be made into a replica with the same homelib as the target.
Displayed when the -6d option is specified. A similar message is displayed for groups.

- WARNING: A differing replica of user *name* exists on the target library, but the replica's homelib (*system*) is not the source library. The source will overwrite the target.

Displayed when the -6o option is specified. A similar message is displayed for groups.

- WARNING: A differing replica of user *name* exists on the target library, but the replica's homelib (*system*) is not the source library. No action will be taken for this user.

Displayed when no options are specified. A similar message is displayed for groups.

- WARNING: Bad internal type for user *name*! No action will be taken for this user.

Internal USGRPMIG program error.

- User's default group *source_default_group* does not exist on target, changing to *target_general_group*.
- User's default storage category *source_default_storage* does not exist on target, changing to *target_system_default_storage*.

A similar message is displayed for groups.

- User's default archive category *source_default_archive* does not exist on target, changing to *target_system_default_archive*.

A similar message is displayed for groups.

- User's added-by user *source_owner* does not exist on target, changing to *target_admin_user*.

A similar message is displayed for groups.

VUPDATE

VUPDATE enters a value in the File Name property for any document's Version object that lacks such a value. IDM Desktop requires a version file name to display icons in search results. If the File Name property has no value, IDM Desktop must do additional time-consuming processing to determine an appropriate substitute value to display, thus degrading search performance. This property value is typically missing in Version objects only for documents checked in by older user interfaces, such as Saros Document Manager.

Usage

```
vupdate propertyserver librarysystem logfile
```

Parameters and Switches

| | |
|-----------------------|--|
| <i>propertyserver</i> | The name of the server where the database and initial library system services are installed for the library system on which you want to run VUPDATE. |
| <i>librarysystem</i> | The name of the library system on which you want to run VUPDATE. |
| <i>logfile</i> | (Optional) The file in which you want VUPDATE to log the information about the documents being updated. |

Usage Notes

- VUPDATE can be run multiple times on the same library system.
- Depending on the amount of data to be updated, VUPDATE can take a long time to run. If you want to perform the update in stages, you can interrupt execution by typing Ctrl-C and resume later by running VUPDATE again.

Procedure

To run VUPDATE:

1. Notify all users to check in any documents they currently have checked out, then close their library system sessions.
2. Make sure the Logons Enabled? property in the System object is set to No.
This allows only one user from the Administrators group to open a library system session.
3. On the initial storage server, access the directory containing the vupdate executable.

| OS Platform | Location |
|-------------|---|
| Windows | <i>StorageManagerInstallDrive:\filenet\LibrarySystem\Util</i> |
| UNIX | <i>StorageManagerInstallDir/filenet/librarysystem/util</i> |

4. On the command line, enter the command to run VUPDATE. For example:

```
vupdate mybox mylib c:\filenet\mylib\vupdate.log
```

5. Set the Logons Enabled? property in the System object to Yes and notify users they can again access the library system.

Creating Reports Using Database Schemas

This topic describes how you can access the database tables for a Content Services 5.5 library system using your SQL report generating software.

Use these database schemas for reporting only! All database table values should be treated as read-only in your reports. Be careful not to update data using these database schemas.

Each documented database table represents an object type in a Content Services (CS) library system. Each documented entry represents a table column corresponding to a particular property of an object. The property description includes the corresponding database field type and length. You may need this field type and length information if you are importing data into another database application which requires knowledge about the type and length of any data values being imported. On the other hand, if you are planning to bring data into a word processing document, you may not need this information at all.

The API Reference information for each entry is found within the parentheses in the Property column of each table. This information reflects the underlying CS Toolkit. This Toolkit is no longer generally available to new customers (and is not needed to generate reports). However, the information included here may be of value to internal FileNet staff (e.g., Professional Services) or long-standing customers who are still maintaining applications built directly on the CS API rather than the IDM Desktop Toolkit.

The last section of this topic has examples of simple SQL queries you can run against a Content Services 5.5 library system.

Changing Existing Reports

Prior to IDM Document Services 4.2, library system objects were exposed using the Report Writer Interface that shipped with the Microsoft SQL Server version of Content Services. The Report Writer Interface feature was then withdrawn from Content Services (i.e., Mezzanine) as of version 4.2. This topic, which documents the Content Services version 5.4 database schemas and the mappings of database tables to Content Services properties, now replaces the former Report Writer Interface and extends the number of objects that you can access and use when generating reports.

If you want to continue to run reports that were created using the now-discontinued Report Writer Interface against Content Services 5.4, you will need to change them to access the database tables listed directly instead of using the Views defined by the Report Writer Interface. The following mappings will help you do this:

- Checkout Object View (CHECKOUTOBJ): From VERSION_CHECKOUT table
- Custom Object View (CUSTOMOBJ): From CUSTOM table
- Group Object View (GROUPOBJ): From GRP table
- Group Member View (GROUPMEMBEROBJ): From GRP_USR table
- Item View (ITEMOBJ): From ELEMENT table, INA_INFO (Custom Properties) table
- Keyword View (KEYWORDOBJ): From ELEM_KW table
- User View (USEROBJ): From USR table
- Version View (VERSIONOBJ): From VERSION table

Database Schema

Tables

- “ACCESS_CONTROL” on page 467
- “APP Table” on page 468
- “ATTR_INFO Table” on page 469
- “AUDIT_EVENTS Table” on page 471
- “AUDIT_LOG Table” on page 472
- “BRANCH_SERVER Table” on page 474
- “CATALOG Table” on page 476
- “CS_SERVER Table” on page 478
- “CSI_INDEX Table” on page 479
- “CSP Table” on page 481
- “CUSTOM Table” on page 482
- “CUSTOM_ATTR_NAME Table” on page 486
- “DEFAULTS Table” on page 487
- “DEFAULT_ACCESS Table” on page 488
- “DISP_STACK Table” on page 489
- “ELEM_KW Table” on page 491
- “ELEMENT Table” on page 492
- “ENAME_CONTROL Table” on page 495
- “Provides client software with connection information needed to connect to a Storage or Content Search Manager.” on page 521
- “GRP Table” on page 497
- “GRP_USR Table” on page 500
- “ICON Table” on page 501
- “INA_INFO Table” on page 502
- “ITEM_MV_CUSTOM01 Table” on page 504
- “ITEM_MV_CUSTOM02 Table” on page 505
- “ITEM_MV_CUSTOM03 Table” on page 506

- “ITEM_MV_CUSTOM04 Table” on page 507
- “ITEM_MV_CUSTOM05 Table” on page 508
- “ITEM_MV_CUSTOM06 Table” on page 509
- “ITEM_MV_CUSTOM07 Table” on page 510
- “ITEM_MV_CUSTOM08 Table” on page 511
- “ITEM_MV_CUSTOM09 Table” on page 512
- “ITEM_MV_CUSTOM10 Table” on page 513
- “LANGID Table” on page 514
- “LIBRARY Table” on page 515
- “MESSAGE Table” on page 520
- “NAME_SERVICE Table” on page 521
- “NUMID Table” on page 522
- “OBJECT_INFO Table” on page 523
- “RELATION Table” on page 525
- “REPL_KW Table” on page 527
- “REPLOBJ Table” on page 528
- “SESSIONTAB Table” on page 533
- “STACK Table” on page 536
- “STORAGE_ALLOC Table” on page 538
- “USR Table” on page 539
- “VCKO_MV_CUSTOM01 Table” on page 543
- “VCKO_MV_CUSTOM02 Table” on page 544
- “VCKO_MV_CUSTOM03 Table” on page 545
- “VCKO_MV_CUSTOM04 Table” on page 546
- “VCKO_MV_CUSTOM05 Table” on page 547
- “VCKO_MV_CUSTOM06 Table” on page 548
- “VCKO_MV_CUSTOM07 Table” on page 549
- “VCKO_MV_CUSTOM08 Table” on page 550
- “VCKO_MV_CUSTOM09 Table” on page 551

- “VCKO_MV_CUSTOM10 Table” on page 552
- “VERS_DISP Table” on page 553
- “VERS_MV_CUSTOM01 Table” on page 554
- “VERS_MV_CUSTOM02 Table” on page 555
- “VERS_MV_CUSTOM03 Table” on page 556
- “VERS_MV_CUSTOM04 Table” on page 557
- “VERS_MV_CUSTOM05 Table” on page 558
- “VERS_MV_CUSTOM06 Table” on page 559
- “VERS_MV_CUSTOM07 Table” on page 560
- “VERS_MV_CUSTOM08 Table” on page 561
- “VERS_MV_CUSTOM09 Table” on page 562
- “VERS_MV_CUSTOM10 Table” on page 563
- “VERSION Table” on page 564
- “VERSION_CHECKOUT Table” on page 568

ACCESS_CONTROL

ACCESS_CONTROL Table (SPI_OBJ_ACCESS_CONTROL, 3840)

Controls the access to a particular User, Group, or Item object by granting access privileges to a specific user or group.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| AC_INST_NM <i>(Parent Object Name)</i> | varchar (208) | varchar2 (208) | Contains the name of the parent User, Group, or Item object to which this Access Control object controls access. |
| AC_U_NAME <i>(User or Group Name)</i> | varchar (32) null | varchar2 (32) | Contains the name of the user or group being granted access privileges. |
| AC_TYPE <i>(User or Group Type)</i> | smallint null | numeric (5) | Indicates the type of object that the name in AC_U_NAME describes. |
| AC_OBJ_TYPE <i>(Parent Object Type)</i> | int null | numeric (10) | Indicates the type of object that the name in AC_INST_NM describes. |
| AC_CODE <i>(Access Level)</i> | smallint null | numeric (5) null | Indicates the access privileges granted to the user or group named in AC_U_NAME. |
| AC_ORDER <i>(Access Order)</i> | tinyint null | numeric (3) null | Internal use only. |

APP Table

APP Table (SPI_OBJ_APPLICATION, 7424)

Obsolete.

ATTR_INFO Table

ATTR_INFO Table (SPI_OBJ_ATTR_INFO, 18944)

Describe the attributes corresponding to all the tables/objects in Content Services property repository.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|-----------------------|------------------------|---|
| AI_ID <i>(Attribute ID)</i> | int | numeric (10) | Contains Attribute ID. |
| AI_NAME <i>(Attribute Name)</i> | varchar (128) null | varchar2 (128) null | Specifies the attribute column name in the corresponding table in database. |
| AI_TABLENAME <i>(Attribute Table Name)</i> | varchar (128) null | varchar2 (128) null | Specifies the table (or object) to which the attribute belongs. |
| AI_TYPE <i>(Attribute Type)</i> | int null | numeric (10) null | Specifies the data type of attribute. |
| AI_LENGTH <i>(Attribute Length)</i> | int null | numeric (10) null | Specifies the size of attribute column in the corresponding table. |

AUDIT_CONTROL Table

AUDIT_CONTROL Table (SPI_OBJ_AUDIT_CONTROL, 4352)

Enables and controls auditing of user events on a library system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| AC_AUDIT_LOG_ON <i>(Logging Enabled?)</i> | tinyint (208) | numeric (3) null | Indicates whether Content Services auditing will be performed for all users of a library system (SPI_YESNO_NO_S or SPI_YESNO_YES_S) |
| AC_CURRENT_SET <i>(Current Set)</i> | int null | numeric (10) | Internal use only. This property value is assigned to each logged audit event in the Audit Log object. The value is incremented every time the stpIncrementSet procedure is run and, under certain conditions, when the stpTruncAuditLog procedure is run. |

AUDIT_EVENTS Table

AUDIT_EVENTS Table (SPI_OBJ_AUDIT_EVENTS, 18176)

Identifies the CS library system audit events that will be logged.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| AE_ID (<i>Event ID</i>) | smallint | numeric(5) | Contains the internal ID of the audit event. |
| AE_LOG_ON_SUCCESS (<i>Log Success?</i>) | tinyint null | numeric(3) null | Indicates whether events that succeed will be logged. |
| AE_LOG_ON_FAILURE (<i>Log Failure?</i>) | tinyint null | numeric(3) null | Indicates whether events that fail will be logged |

AUDIT_LOG Table

AUDIT_LOG Table (SPI_OBJ_AUDIT_LOG, 18432)

Describes a logged audit event.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| AL_EVENT_ID (Event ID) | smallint | numeric(5) null | Contains the event ID of the audit event. |
| AL_SEQUENCE (Event Sequence) | int null | numeric(10) | Internal use only. |
| AL_OBJECT_KEY (Object Key) | varchar(165) null | varchar2(165) null | Contains the key value of the audit event. This property only applies to events that require a key value. |
| AL_EVENT_PARAM1 (Event Param1) | varchar(254) null | varchar2(254) null | Contains the input/output audit event parameter values. Data that exceeds the length of this property will be truncated. |
| AL_EVENT_PARAM2 (Event Param2) | varchar(254) null | varchar2(254) null | Contains the input/output audit event parameter values. Data that exceeds the length of this property will be truncated. |
| AL_EVENT_PARAM3 (Event Param3) | varchar(254) null | numeric(5) | Contains the input/output audit event parameter values. Data that exceeds the length of this property will be truncated. |
| AL_EVENT_PARAM4 (Event Param4) | varchar(254) null | varchar2(254) null | Contains the input/output audit event parameter values. Data that exceeds the length of this property will be truncated. |
| AL_SR_START_KEY (Start Key) | varchar(254) null | varchar2(254) null | Contains the start key for a search when the search is continued (that is, indicates a starting point for the results set). Applies only to the spiQuerySearchRequest, spiQueryList, and spiQueryCard audit events. |
| AL_SR_DIRECTION (Direction) | int null | numeric(10) null | Contains the search direction (ascending or descending) indicated for the results set. Applies only to the spiQuerySearchRequest, spiQueryList, and spiQueryCard audit events. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|------------------------------|-------------------------------|---|
| AL_SR_COUNT <i>(Count)</i> | int null | numeric(10) null | Contains the number of hits returned by the search (or NULL if the search fails). Applies only to the spiQuerySearchRequest, spiQueryList, and spiQueryCard audit events. |
| AL_DATETIME <i>(Date Time)</i> | datetime null | date | Contains the date and time the audit event occurred. |
| AL_USER <i>(User Name)</i> | varchar(32) null | varchar2(32) | Contains the name of the library system user who initiated the audit event. |
| AL_STATUS <i>(Status)</i> | int null | numeric(10) null | Contains the status code returned by the audit event. |
| AL_WORKSTATN_ADDR <i>(Workstation)</i> | varchar(32) null | varchar2(32) null | Contains the workstation address of the client where the audit event was originated. |
| AL_ELAPSED_TIME <i>(Elapsed Time)</i> | int null | numeric(10) null | Contains the elapsed time (in milliseconds) that was required to complete the audit event. |
| AL_SET_NUMBER <i>(Set Number)</i> | int null | numeric(10) null | Contains the set number from the Audit Control Object Current Set property (AC_CURRENT_SET). |
| AL_PROCESSID <i>(Process ID)</i> | int null | numeric(10) | Contains the process ID, assigned by the operating system, of the program that initiated the audit event. |
| AL_FLAGS <i>(Flags)</i> | int null | numeric(10) null | Contains the flags specified for the audit event, if applicable. |
| AL_FILENAME <i>(File Name)</i> | varchar(254) null | varchar2(254) null | Contains the file name referenced by the audit event, if applicable. |

BRANCH_SERVER Table

BRANCH_SERVER Table (SPI_OBJ_SHELF_SERVER, 3328)

Describes a Storage Manager, which not only controls storage repositories used to store system-managed files but also services operational requests.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| BS_NODE (<i>Server</i>) | varchar(32) | varchar2(32) | Contains the unique network name of the server on which the Storage Manager is located. |
| BS_NAME (<i>Service Name</i>) | varchar(32) null | varchar2(32) null | Contains the unique name assigned to the Storage Manager during library system server installation. |
| BS_CLASS (<i>Service Class</i>) | varchar(32) null | varchar2(32) null | (Internal use only) Contains the class (for example, Windows) of the server on which the Storage Manager is located. |
| BS_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about the Storage Manager. |
| BS_START_DATE (<i>Start Date</i>) | datetime null | date null | Indicates the date and time that the Storage Manager was most recently started. |
| BS_STOP_DATE (<i>Stop Date</i>) | datetime null | date null | Indicates the date and time that the Storage Manager was most recently stopped. |
| BS_STATUS (<i>Status</i>) | int null | numeric(10) null | Indicates whether this Storage Manager is currently operational and/or is available for use. |
| BS_LAST_BACKUP (<i>Backup Date</i>) | datetime null | date null | Contains the date and time of the most recent system backup. |
| BS_HEARTBEAT (<i>Last Heartbeat</i>) | datetime null | date null | Contains the date and time the reclaim daemon for the Storage Manager last attempted to reclaim version files from the associated reclaim repositories. Because this value is updated at regular intervals, it can indicate whether or not the Storage Manager is still active. |
| BS_THING_1 (<i>Future</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|------------------------------|-------------------------------|---|
| BS_THING_2 <i>(Future)</i> | varchar(32) null | varchar2(32) null | Reserved for future use. |
| BS_VERSION <i>(Software Version)</i> | varchar(32) null | varchar2(32) null | Contains the software version number of the Storage Manager. |
| BS_INSTALLED_DRIVE <i>(Device)</i> | varchar(32) null | varchar2(32) | Contains the device (for example, drive) on which the Storage Manager is installed. |
| BS_ENDIAN <i>(Endian)</i> | int null | numeric(10) null | (Internal use only) Indicates the byte order of a CPU architecture. |
| BS_NODE_ADDR <i>(Server Address)</i> | varchar(32) null | varchar2(32) null | Contains the address of the server where the Storage Manager is located. This property is primarily useful for UNIX systems. |
| BS_SERVICE_ADDR <i>(Service Address)</i> | varchar(32) null | varchar2(32) null | Contains the service port address of the server where the Storage Manager is located (such as, BRSERVER). This property is primarily useful for UNIX systems. |

CATALOG Table

CATALOG Table (SPI_OBJ_CATALOG, 6400)

Describes a Property Manager. The system automatically creates a Property Manager object when Content Services is installed.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| C_NAME (<i>System Name</i>) | varchar(32) | varchar2(32) | Contains the unique name assigned to the Property Manager during server setup. |
| C_NODE (<i>Server</i>) | varchar(32) null | varchar2(32) null | Contains the unique name identifying the network server on which the Property Manager is located. |
| C_STATUS (<i>Status</i>) | int null | numeric(3) null | Indicates whether the Property Manager is currently running. |
| C_START_DATE (<i>Start Date</i>) | datetime null | date null | Indicates when the Property Manager was last started. |
| C_STOP_DATE (<i>Stop Date</i>) | datetime null | date null | Indicates when the Property Manager was last shut down, making it unavailable to users. |
| C_DATA_SPACE (<i>Repository Max Size</i>) | int null | numeric(10) null | Indicates the total amount of disk space assigned to the Property Manager for data storage. |
| C_LOG_SPACE (<i>Transaction Log Max Size</i>) | int null | numeric(10) null | No longer used. |
| C_DATA_PCT (<i>Repository Space Used (%)</i>) | int null | numeric(10) null | Indicates the percentage of Property Manager data space that has been used. |
| C_LOG_PCT (<i>Transaction Log Space Used (%)</i>) | int null | numeric(10) null | On library systems that use a Microsoft SQL Server database engine, indicates the percentage of Property Manager transaction log space that has been used. This property does not apply to library systems that use an Oracle database engine. |
| C_THING_1 (<i>Future</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |
| C_THING_2 (<i>Future</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |
| C_VERSION (<i>DB Engine Software Version</i>) | varchar(32) null | varchar2(32) null | Indicates the software version of the Property Manager. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| C_LOG_WARNING <i>(Log Warning Threshold (%))</i> | int null | numeric(10) null | On library systems that use a Microsoft SQL Server database engine, sets the Property Manager transaction log warning threshold. This property does not apply to library systems that use an Oracle database engine. |
| C_LOG_FULL <i>(Log Full Threshold (%))</i> | int null | numeric(10) null | On library systems that use a Microsoft SQL Server database engine, sets the Property Manager transaction log full threshold. This property does not apply to library systems that use an Oracle database engine. |
| C_DATA_WARNING <i>(Repository Warn Threshold (%))</i> | int null | numeric(10) null | Sets the Property Manager warning threshold value for repositories. |
| C_DATA_FULL <i>(Repository Full Threshold (%))</i> | int null | numeric(10) null | Sets the Property Manager full threshold value for repositories. |
| C_SHARED <i>(Customer Supplied DB Engine?)</i> | int null | numeric(10) null | Indicates whether the database engine is being shared with other systems or applications. |
| C_INSTALL_DRIVE <i>(Device)</i> | varchar(32) null | varchar2(32) null | Identifies the storage device on which the Property Manager for this system was installed. |
| C_NODE_ADDR <i>(Server Address)</i> | varchar(32) null | varchar2(32) null | Identifies the network address for the server on which the Property Manager for this system is installed. |
| C_SERVICE_ADDR <i>(Service Address)</i> | varchar(32) null | varchar2(32) null | Contains the service address for the server on which the Property Manager is installed. |
| C_PROP_INDX_TBSP | varchar(32) null | varchar(32) null | Contains the name of the tablespace that stores indexes of the properties of the library system. (Oracle databases only) |

CS_SERVER Table

CS_SERVER Table (SPI_OBJ_CS_SERVER, 9216)

Describes a Content Search Manager that contains repositories for content searching.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| CSS_NAME (Service Name) | varchar(32) | varchar2(32) null | Contains the name of this Content Search Manager. |
| CSS_NODE (Server) | varchar(32) null | varchar2(32) | Contains the name of the machine on which this Content Search Manager is located. |
| CSS_DEVICE (Device) | varchar(32) null | varchar2(32) null | Contains the name of the top level device (for example, the drive letter) where this Content Search Manager is located. |
| CSS_PATH (Path) | varchar(96) null | varchar2(96) null | Contains the full path of the root directory containing this Content Search Manager. |
| CSS_COMMENT (Comment) | varchar(224) null | varchar2(224) null | Contains a comment string. |
| CSS_START_DATE (Start Date) | datetime null | date null | Contains the date and time that this Content Search Manager was started. |
| CSS_STOP_DATE (Stop Date) | datetime null | date null | Contains the date and time that this Content Search Manager was stopped. |
| CSS_STATUS (Status) | int null | numeric(10) null | Indicates the operational status of this Content Search Manager. |
| CSS_VERSION (Software Version) | varchar(32) null | varchar2(32) null | Contains the software version of this Content Search Manager. |
| CSS_MAX_CONNS (Max Number of Connections) | int null | numeric(10) null | Internal use only. |
| CSS_NODE_ADDR (Server Address) | varchar(32) null | varchar2(32) null | Contains the server address for this Content Search Manager. |
| CSS_SERVICE_ADDR (Service Address) | varchar(32) null | varchar2(32) null | Contains the service address for this Content Search Manager. |
| CSS_VERITY_LOCALE (Verity Locale) | varchar(32) null | varchar(32) null | Contains the Verity locale for this Content Search Manager. |

CSI_INDEX Table

CSI_INDEX Table (SPI_OBJ_CSI_INDEX, 9472)

Describes a repository for a Content Search Manager.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|---------------------|----------------------|--|
| CSI_ID <i>(Content Search Repository ID)</i> | int | numeric(10) | Contains the system supplied ID of this content search repository (for example, 1). |
| CSI_NAME <i>(Content Search Repository Name)</i> | varchar(32) null | varchar2(32) null | Contains the system supplied name of this content search repository (for example, index001). |
| CSI_NODE <i>(Server)</i> | varchar(32) null | varchar2(32) null | Contains the unique name identifying the network server on which this content search repository is located. |
| CSI_DEVICE <i>(Device)</i> | varchar(32) null | varchar2(32) null | Contains the name of the top level device (for example, the drive letter) where this content search repository is located. |
| CSI_PATH <i>(Path)</i> | varchar(96) null | varchar2(96) null | Contains the full path of the root directory containing this content search repository. |
| CSI_MERGE_HOUR <i>(Merge Time of Day)</i> | int null | numeric(10) null | Indicates the time of day in 24 hour format that the CSMERGE utility is run. |
| CSI_MAX_SIZE <i>(Max Repository Size)</i> | int(4) null | int(4) null | Indicates the maximum size (bytes) of the parts directory of the collection (index). |
| CSI_ERRCOUNT <i>(Number of Indexing Errors)</i> | int null | numeric(10) null | Indicates how many errors were detected the last time the indexing daemon ran. |
| CSI_PRIORITY <i>(Indexing Priority)</i> | int null | numeric(10) null | Reserved for future use. |
| CSI_SEMA <i>(Indexing Semaphore)</i> | smallint null | numeric(5) null | Internal use only. |
| CSI_STATUS <i>(Indexing Status)</i> | smallint null | numeric(5) null | Indicates the operational status of this content search repository. |
| CSI_LASTMERGE_TIME <i>(Last Merge Time)</i> | datetime null | date null | Internal use only. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|------------------------------|-------------------------------|---|
| CSI_STOP_TIME <i>(Indexing Stop Time)</i> | datetime null | date null | Internal use only. |
| CSI_CSMERGE_PID <i>(Merging PID)</i> | smallint null | numeric(10) null | Internal use only. |
| CSI_CSINDEX_PID <i>(Indexing PID)</i> | smallint null | numeric(10) null | Internal use only. |
| CSI_WORKFILE_SIZE <i>(Working File Size)</i> | int null | numeric(10) null | Specifies the maximum size in bytes of a temporary sort file. |
| CSI_WORKMEM_SIZE <i>(Working File Size)</i> | int null | numeric(10) null | Specifies the size in bytes of a temporary buffer that is allocated. |
| CSI_LOGSQL <i>(SQL Logging)</i> | tinyint null | numeric(3) null | Specifies whether the SQL commands used by the CSINDEX and CSMERGE utilities will be logged to the csi00n.log file. |
| CSI_RMAJOR <i>(Full Backup ID)</i> | smallint null | numeric(5) null | Internal use only. |
| CSI_MAX_IMMED_SIZE <i>(Max Immediate File Size)</i> | int null | numeric(10) null | Specifies the maximum size (in bytes) that the immediate index file can reach before the CSMERGE utility is run |
| CSI_NET_PATH <i>(Network Path)</i> | varchar(254) null | varchar2(254) null | Contains the network path location of a content search repository located on a network server. |
| CSI_OVERRIDE <i>(Override Defaults)</i> | tinyint null | numeric(3) null | Contains the override defaults. |

CSP Table

CSP Table

Stores the indexing/de-indexing requests for a document added to Content Services.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| CSP_CSI_ID (<i>CSI ID</i>) | int null | numeric (10) null | Content Search Repository Identifier. |
| CSP_CSI_VERSION (<i>Version</i>) | varchar (7) null | char (7) null | Specifies the version of item to be indexed/de-indexed. |
| CSP_CSI_ITEMID (<i>Item ID</i>) | varchar (9) null | char (9) null | Specifies Item ID of item to be indexed/de-indexed. |
| CSP_CSI_ACTION (<i>CSI Action</i>) | tinyint null | numeric (3) null | Specify the type of operation. |
| CSP_CSI_CID (<i>CSI CID</i>) | varchar (32) null | varchar2 (32) null | Contains the unique ID that identifies the current version of the item in the content search repository. |
| CSP_CSI_USER (<i>CSI User</i>) | varchar (32) null | varchar2 (32) null | Contains the user that added the document in queue. |
| CSP_CSI_DATE_ADDED (<i>Date Added</i>) | datetime null | date null | Specifies the date of adding document in queue. |
| CSP_CSI_FLIST (<i>CSI Flist</i>) | varchar (32) null | varchar2 (32) null | For internal use. |
| CSP_CSI_FILE_SIZE (<i>File Size</i>) | int null | numeric (10) null | Specifies size of file in bytes. |

CUSTOM Table

CUSTOM Table (SPI_OBJ_CUSTOM_OBJECT, 8960)

Describes a Custom object that is used to extend the FileNet Content Services API in your system. This object may be either a high level object (for example, an Enhanced User object) or a lower level object that will function as a quasi list (for example, a Year End Evaluation object under the User object).

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| CO_ISV_ID (<i>ISV ID</i>) | varchar(8) | varchar2(8) | Contains a Content Services assigned string identifier for the object. |
| CO_OBJECT_TYPE (<i>Object Type</i>) | varchar(32) null | varchar2(32) | Contains a developer assigned string identifier for the object. |
| CO_KEY (<i>Object Key</i>) | varchar(165) null | varchar2(165) | Contains a unique, developer assigned key to the object instance. |
| CO_STRING1 (<i>String1</i>) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING2 (<i>String2</i>) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING3 (<i>String3</i>) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING4 (<i>String4</i>) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING5 (<i>String5</i>) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING6 (<i>String6</i>) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--------------------------------------|----------------------|-----------------------|--|
| CO_STRING7 (String7) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING8 (String8) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING9 (String9) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_STRING10 (String10) | varchar(32) null | varchar2(32) null | Contains one of ten possible customized string properties (each up to 32 characters) that describe the object. |
| CO_LSTRING1 (Long String1) | varchar(254) null | varchar2(254) null | Contains one of five possible longer, customized string properties (each up to 254 characters) that describe the object. |
| CO_LSTRING2 (Long String2) | varchar(254) null | varchar2(254) null | Contains one of five possible longer, customized string properties (each up to 254 characters) that describe the object. |
| CO_LSTRING3 (Long String3) | varchar(254) null | varchar2(254) null | Contains one of five possible longer, customized string properties (each up to 254 characters) that describe the object. |
| CO_LSTRING4 (Long String4) | varchar(254) null | varchar2(254) null | Contains one of five possible longer, customized string properties (each up to 254 characters) that describe the object. |
| CO_LSTRING5 (Long String5) | varchar(254) null | varchar2(254) null | Contains one of five possible longer, customized string properties (each up to 254 characters) that describe the object. |
| CO_INTEGER1 (Integer1) | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|-------------------|---------------------|--|
| CO_INTEGER2 <i>(Integer2)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER3 <i>(Integer3)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER4 <i>(Integer4)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER5 <i>(Integer5)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER6 <i>(Integer6)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER7 <i>(Integer7)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER8 <i>(Integer8)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER9 <i>(Integer9)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_INTEGER10 <i>(Integer10)</i> | int null | numeric(10) null | Contains one of ten possible customized integer properties (each up to 9 characters) that describe the object. |
| CO_DATETIME1 <i>(DateTime1)</i> | datetime null | date null | Contains one of five possible customized date/time properties that describe the object. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|------------------------------|-------------------------------|--|
| CO_DATETIME2 <i>(DateTime2)</i> | datetime null | date null | Contains one of five possible customized date/time properties that describe the object. |
| CO_DATETIME3 <i>(DateTime3)</i> | datetime null | date null | Contains one of five possible customized date/time properties that describe the object. |
| CO_DATETIME4 <i>(DateTime4)</i> | datetime null | date null | Contains one of five possible customized date/time properties that describe the object. |
| CO_DATETIME5 <i>(DateTime5)</i> | datetime null | date null | Contains one of five possible customized date/time properties that describe the object. |
| CO_VERSION1 <i>(Version1)</i> | varchar(7) null | char(7) null | Contains one of three possible customized version properties to describe the object. |
| CO_VERSION2 <i>(Version2)</i> | varchar(7) null | char(7) null | Contains one of three possible customized version properties to describe the object. |
| CO_VERSION3 <i>(Version3)</i> | varchar(7) null | char(7) null | Contains one of three possible customized version properties to describe the object. |
| CO_TEXT1 <i>(Text 1)</i> | text null | long null | Contains a memo property of 2048 maximum bytes. The CO_LSTRING1 property has been lengthened to 2048 bytes to provide similar functionality. |
| CO_OWNER <i>(Added By User)</i> | varchar(32) null | varchar2(32) null | Contains the name of the user who added the Custom object. |
| CO_DATE <i>(Date Added)</i> | datetime null | date null | Contains the date the Custom object was added. |

CUSTOM_ATTR_NAME Table

CUSTOM_ATTR_NAME Table (SPI_OBJ_CUSTOM_ATTR_NAME, 5632)

This object is typically used for national language support. It contains a custom property display string (for example, for a particular language). When you add a custom property via CS Admin Tools, an instance of this object is added with a Display Name value to match the Property Name value in the new Custom Property object. An instance of this property must be added for each language string that is to be displayed. A new Display Names property in each Custom Property object lists these various Custom Property Name objects that might be created. Only administrators can access this object and change the Display Name value.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| CAN_CA_ID (<i>Custom Property ID</i>) | int | numeric(10) | Uniquely identifies the custom property. |
| CAN_LANG_ID (<i>Display Language ID</i>) | tinyint null | numeric(3) | Identifies the display language to be used. |
| CAN_NAME (<i>Display Name</i>) | varchar(32) null | varchar2(32) null | Contains the string to be displayed to users as the name of this custom property (typically, in conjunction with a particular language). |
| CAN_CA_PARENT_OBJ (<i>Parent Object Number</i>) | int null | numeric(10) | Valid parent object numbers are 256 (Item), 512 (Version), and 768 (Checkout). |
| CAN_CA_MUL_VALUED (<i>Multi-valued?</i>) | tinyint null | numeric(3) | Indicates whether this custom property is multi valued or single valued. |

DEFAULTS Table

Describes the system default values for various attributes.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|-----------------------|------------------------|---|
| DEFAULT_VALUE <i>(Default Value)</i> | varchar (254) null | varchar2 (254) null | Specifies default value to be used for an attribute. |
| DEFAULT_KEY <i>(Default Key)</i> | int | numeric (10) | Specifies attribute key mapped to a particular attribute. |

DEFAULT_ACCESS Table

DEFAULT_ACCESS Table (SPI_OBJ_DEFAULT_ACCESS, 6912)

Causes Content Services to automatically enter an Access Control object that will control a user or group access to a newly created Itemobject at the time the new object is added to the property repository.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| DA_OWNER_TYPE <i>(Source Object Type)</i> | int | numeric(10) | Contains the type of parent object (SPI_OBJ_SYSTEM, SPI_OBJ_USER, or SPI_OBJ_GROUP) that the name in the DA_OWNER_NAME property describes. |
| DA_OWNER_NAME <i>(Source Object Name)</i> | varchar(32) null | varchar2(32) | Contains the object type (SPI_OBJ_USER or SPI_OBJ_GROUP) to be placed into the AC_TYPE property in the SPI_OBJ_ACCESS_CONTROL object. |
| DA HOLDER_TYPE <i>(User or Group Type)</i> | smallint null | numeric(5) | Contains the string to be displayed to users as the name of this custom property (typically, in conjunction with a particular language). |
| DA HOLDER_NAME <i>(User or Group Name)</i> | varchar(32) null | varchar2(32) | Contains the name of the user or group to be placed into the AC_U_NAME property in the SPI_OBJ_ACCESS_CONTROL object. |
| DA_OBJECT_TYPE <i>(Object Type)</i> | int null | numeric(10) | Indicates the type of object that the SPI_OBJ_ACCESS_CONTROL object will control access to. (Currently, only SPI_OBJ_ITEM is supported.) |
| DA_CODE <i>(Access Level)</i> | smallint null | numeric(5)null | Contains the access privileges that will be placed into the AC_CODE property in the SPI_OBJ_ACCESS_CONTROL object. |

DISP_STACK Table

DISP_STACK Table (SPI_OBJ_DISP_SHELF, 3072)

Describes an archive repository that has been added to the library system to be used in the version archive process as a temporary storage directory for versions being moved offline.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| DS_ID (<i>Archive Repository ID</i>) | int | numeric(10) | Contains the unique identifier for the archive repository. |
| DS_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about the archive repository. |
| DS_RECLAIMABLE (<i>Reclaimable?</i>) | int null | numeric(10) null | Indicates whether versions that are moved offline using the archive categories linked to the archive repositories can be reclaimed. |
| DS_NODE (<i>Server</i>) | varchar(32) null | varchar2(32) null | Contains the unique name of the network server on which this archive repository is located. |
| DS_DEV (<i>Device</i>) | varchar(32) null | varchar2(32) null | Contains the name of the storage device where the archive repository is located. |
| DS_DIR (<i>Directory</i>) | varchar(32) null | varchar2(32) null | Contains the name of the directory used to locate the archive repository relative to the server and the device name. |
| DS_SERVR_NM (<i>Server Name</i>) | varchar(32) null | varchar2(32) null | Contains the constant value BRSERVER. Reserved for future use. |
| DS_SERVR_CLASS (<i>Server Class</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |
| DS_PATH (<i>Path</i>) | varchar(96) null | varchar2(96) null | Contains the path location of the archive repository. |
| DS_THING_1 (<i>Future</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |
| DS_THING_2 (<i>Future</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |
| DS_DELETE_FILES (<i>Delete Upon Archive?</i>) | tinyint null | numeric(3) null | Indicates whether to forego copying archived version files to the archive repository and delete them immediately instead. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| DS_NET_PATH <i>(Network Path)</i> | varchar(254) null | varchar2(254) null | Contains the network path location of an archive repository located on a network server. |

ELEM_KW Table

ELEM_KW Table (SPI_OBJ_KEYWORD, 4096)

Provides information about a keyword that is associated with a particular item. Keyword objects appear as child objects of an Item object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--------------------------------------|----------------------|-----------------------|--|
| KW_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Contains the Item ID of the item the keyword is associated with. |
| KW_NAME <i>(Keyword)</i> | varchar(224) null | varchar2(224) null | Contains the user supplied keyword value. |

ELEMENT Table

ELEMENT Table (SPI_OBJ_ITEM, 256)

Describes a protected or unprotected item that has been added to the library system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|------------------------|--|
| E_NAME (<i>Item ID</i>) | varchar(9) | char(9) | Contains the unique numeric identifier for the item. |
| E_OWNER (<i>Added By User</i>) | varchar(32) null | varchar2(32) null | Contains the name of the user who added the Item object to the property repository. |
| E_GROUP (<i>Added By Group</i>) | varchar(32) null | varchar2(32) null | Contains the name of the active group of the user who added the Item object to the property repository. |
| E_CREATE_DATE (<i>Date Added</i>) | datetime null | date null | Contains the date and time the Item object was added to the property repository. |
| E_COMMENT (<i>Comment</i>) | text null | varchar2(2000) null | Contains a comment about the item. |
| E_PROTECTED (<i>Protected?</i>) | tinyint null | numeric(3) null | Indicates whether the item is a protected item. |
| E_ORG_FILENAME (<i>Original File Name</i>) | varchar(32) null | varchar2(256) null | Contains the file name for the first version of the item. The Oracle value also contains the file name extension. Note that this is a Text data type, and cannot be indexed. |
| E_ORG_FILENAME_EXT (<i>File Name Extension</i>) | text null | varchar2(2000) null | Contains the extension of the file name for the first version of the item. For Oracle, this value is contained in E_ORG_FILENAME. Note that this is a Text data type, and cannot be indexed. |
| E_ALLOC_TYPE (<i>Storage Category</i>) | varchar(32) null | varchar2(32) null | (Protected items only) Contains the name of the storage category that indicates where (in which storage repository) checked in version files for the item will be stored. |
| E_CHKOUT_FILE (<i>Obsolete.</i>) | Obsolete. | Obsolete. | Obsolete. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|-------------------|--------------------|--|
| E_VARIANTS (Allow Secondary Version Lines?) | tinyint null | numeric(3) null | Controls whether secondary version lines are allowed for the item. |
| E_ONLINE_LMT (Online Limit) | int null | numeric(10) null | (Protected items only) Establishes the maximum number of versions of the item that can be stored online before automatic archiving will begin moving the oldest version offline. |
| E_DISP_MODE (Archive Category) | varchar(32) null | varchar2(32) null | (Protected items only) Contains the name of the archive category that indicates where (to which archive repository) versions will be moved during the automatic archival process. |
| E_CURVER_NUM (Current Version ID) | varchar(7) null | varchar2(7) null | Contains the unique identifier for the current version of the item. |
| E_CURVER_DATE (Current Version Date) | datetime null | date null | Contains the date and time the Item's current version was added to the library system. |
| E_CURVER_AUTHOR (Current Version Author) | varchar(32) null | varchar2(32) null | Contains the name of the user who checked in the current version of the item. |
| E_ONLINE_CNT (Online Version Count) | int null | numeric(10) null | (Protected items only) Contains the total number of versions currently online in Content Services for this item. |
| E_ONLINE_SIZE (Online Disk Space) | int null | numeric(10) null | (Protected items only) Indicates the total number of characters (bytes) of online storage occupied by the versions of the item. The maximum online storage size is 2 GB. In case this value exceeds then no new version of an item can be checked-in until some older versions are archived. |
| E_LAST_DATE (Last Access Date) | datetime null | date null | Contains the date and time any version of this item was last accessed (when a user last checked out, checked in, or copied a version). |
| E_LAST_VER (Last Accessed Version) | varchar(7) null | varchar2(7) null | Contains the version ID of the version that was most recently accessed (checked out, checked in, or copied). |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| E_CURVER_CKO <i>(Current Version Checked Out?)</i> | tinyint null | numeric(3) null | Indicates whether the current version is checked out. |
| E_ICON_TITLE <i>(Title)</i> | varchar(254) null | varchar2(254) null | Contains the name for the Item desktop icon displayed by the CS Admin Tools user interface. |
| E_CSI_ID <i>(Current Content Search Repository)</i> | int null | numeric(10) null | Contains the unique ID assigned to the content search repository in which the current version of this item is indexed. |
| E_CSI_CID <i>(Current Version Index ID)</i> | int null | numeric(10) null | Contains the unique ID that identifies the current version of this item in the content search repository. |
| E_CSINDEX_DEFAULT <i>(Default Index Versions?)</i> | tinyint null | numeric(3) null | Indicates whether new versions of the item will be indexed by default. |
| E_CSINDEX_COUNT <i>(Number of Indexed Versions)</i> | smallint null | numeric(5) null | Contains an integer that indicates how many versions of the item are indexed in the content search repository. |
| E_CSINDEX_LIMIT <i>(Indexed Version Limit)</i> | smallint null | numeric(5) null | Indicates whether to override the maximum DICTIONARY file size limit for content search indexing. The default value is No. Set this value to Yes only if (1) the content search repository reaches or nears the maximum DICTIONARY file size limit and (2) only after you have created a new content search repository and reassigned to it all the storage categories that were assigned to the old content search repository. This action will allow deindexing to continue on the old repository while all subsequent indexing occurs on the new repository. |
| E_MIN_DEL_ACCESS <i>(Min Item Delete Access Level)</i> | smallint null | numeric(5) null | Indicates what level of access rights is required to delete Item objects. Regardless of this setting, if an item has no associated versions, anyone with Owner access rights to the Item object can delete it. |
| E_ISREPLICA <i>(Replica Item?)</i> | tinyint null | numeric(3) null | Indicates whether or not the item is a replica item. |

ENAME_CONTROL Table

Describes the document identification number information.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--------------------------------------|----------------------|-----------------------|--|
| C_DATE (Date) | null | date null | Specifies current date to control the generation of sequential Item IDs. |
| E_NUM (Attribute Name) | int null | numeric (10) null | Used to reduce deadlock contention. |
| EC_SEQNO (Sequence Number) | int null | numeric (10) null | Specifies the sequence number to be used to generate the next Item ID. |
| EC_SEQMODE (Sequence Mode) | int null | numeric (10) null | Controls use of sequential Item IDs. |

FLDR_KW Table

FLDR_KW Table (SPI_OBJ_KEYWORD, 4096)

Provides information about a keyword that is associated with a particular item. Keyword objects appear as child objects of an Item object. Provides the information about the items and their containing folders.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| KW_E_NAME (<i>Item ID</i>) | varchar (9) | char (9) | Specifies Item ID of the item stored in a folder. |
| KW_NAME (<i>Folder ID</i>) | varchar (244) | varchar2 (244) | Folder relationship keywords begin with the marker '{MTISDM}'. This marker is followed by the folder id of the folder where the item is filed in. |

GRP Table

GRP Table (SPI_OBJ_GROUP, 1536)

Describes a group that has been added to the library system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| G_NAME (Group Name) | varchar(32) | varchar2(32) | Contains the unique name of the group. |
| G_OWNER_NAME (Added By User) | varchar(32) null | varchar2(32) null | Contains the name of the user who added the Group object to the property repository. |
| G_CREATE_DATE (Date Added) | datetime null | date null | Contains the date and time when the Group object was added to the property repository. |
| G_COMMENT (Comment) | varchar(224) null | varchar2(224) null | Contains a comment about the group. |
| G_STOR_ALLOC (Default Storage Category) | varchar(32) null | varchar2(32) null | Contains the value to apply to the SPI_SES_DEF_FILE_CATEGORY property in the Session object while the session is open and active. |
| G_DISPOSITION (Default Archive Category) | varchar(32) null | varchar2(32) null | Provides a value for the SPI_SES_DEF_DISP_CATEGORY property in the Session object while the session is open and active. |
| G_CHKOUT_F_NAME (Obsolete) | Obsolete. | Obsolete. | Obsolete. |
| G_CHKOUT_F_EXT (Default Checkout File Ext) | varchar(3) null | varchar2(3) null | No longer used. |
| G_PRIMARY_SERVER (Primary Storage Manager) | varchar(32) null | varchar2(32) null | Provides a value for the SPI_SES_PRIMARY_SERVER property in the Session object while the session is open and active. The value indicates the Storage Manager that Content Services will attempt to use first to handle an operational request from a member of this group. |
| G_STATUS (Enabled?) | int null | numeric(10) null | Indicates whether a group is active. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| G_CSINDEX_DEFAULT <i>(Default Index Versions?)</i> | tinyint null | tinyint null | Indicates a default value for the Default Index Versions? property in Item objects. |
| G_SEARCH_DOMAIN <i>(Access Domain)</i> | int null | numeric(10) null | <p>Contains one of the following group domain values indicating the group affiliation(s) to apply when determining which search results to return during an active group member search:</p> <p>SPI_ACCDOMAIN_DEFAULT_S</p> <p>Search results based on the access domain value passed from the current System user object.</p> <p>SPI_ACCDOMAIN_ACTIVEGROUP_S</p> <p>Search results based on access rights assigned to the active user group.</p> <p>SPI_ACCDOMAIN_ALLGROUPS_S</p> <p>Search results based on based on the highest access rights assigned to any of the groups in which the user is a member.</p> |
| G_LANGUAGE <i>(Display Language)</i> | tinyint null | numeric(3) null | Indicates the default language to use for the group. |
| G_SECURE_SEARCH <i>(Secure Searches?)</i> | tinyint null | numeric(3) null | Indicates whether members of this group must have at least Viewer access rights to an item to have that item included in the results of their searches. |
| G_EMAIL_ADDRESS <i>(Email Address)</i> | varchar(64) null | varchar2(64) null | Contains the email address for the active group members. |
| G_CS_THRESHOLD <i>(Content Srch Threshold)</i> | int null | numeric(10) null | Determines whether to do a content search or property search first, when both types of criteria are specified for a search. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| G_CS_DIR_SYNC <i>(Imported Group Type)</i> | int null | numeric(10) null | Enables synchronization of CS library group names with Windows NT domain(s) group names. The User Import tool (dirsync.exe) runs on the initial storage server as a daemon under a selected library system's Storage Manager. When enabled, the daemon performs ongoing one-way synchronization of group and user information from selected network domains to a selected library system database; it does not propagate changes from the library system database back to the network domains. |

GRP_USR Table

GRP_USR Table (SPI_OBJ_GROUP_USER, 4608)

Describes the members of existing CS groups in the library system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| GU_G_NAME <i>(Group Name)</i> | varchar(32) | varchar2(32) | Contains the name of the group to which the user named in the GRPMEM_USER property belongs. |
| GU_U_NAME <i>(User Name)</i> | varchar(32) null | varchar2(32) | Contains the user name of the person who is the group member. |

ICON Table

ICON Table (SPI_OBJ_ICON, 5888)

Describes a user-interface graphical icon associated with a particular user (for example, a Checkout icon or a Search icon).

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| I_U_NAME <i>(User Name)</i> | varchar(32) | varchar2(32) | Contains the user name of the person with which the icon is associated. |
| I_OBJ_TYPE <i>(Object Type)</i> | int null | numeric(10) | Contains the type of object that the icon represents. |
| I_OBJ_KEY <i>(Object Key)</i> | varchar(76) null | varchar2(76) | Contains the key to the object that the icon represents. |
| I_OBJ_TITLE <i>(Icon Title)</i> | varchar(32) null | varchar2(32) | Contains the name that will identify the icon when it is displayed. |

INA_INFO Table

INA_INFO Table (SPI_OBJ_CUSTOM_ATTRIBUTE, 3584)

Describes a custom property. Once created, a custom property cannot be deleted.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|------------------------|---|
| INA_NAME (Property Name) | varchar(32) | varchar2(32) | Contains the name of the property. Unlike the value of INA_ID, this value can be modified. Typically, this value is used as the user interface display name for the custom property. |
| INA_TYPE (Property Type) | int null | numeric(10) null | Indicates the type of data that can be placed in the custom property. |
| INA_LENGTH (Property Length) | int null | numeric(10) null | Indicates the maximum length for the value of the custom property. |
| INA_COMMENT (Comment) | text null | varchar2(2000) null | Contains a comment about the custom property. |
| INA_ID (Property ID) | int null | numeric(10) | Contains the custom property ID. |
| INA_NO_ACCESS (No Privileges) | int null | numeric(10) | Is not available because a user with access rights of none to the Item object in which the custom property appears will always have privileges of none to the value of the custom property. |
| INA_VIEW_ACCESS (Viewer Privileges) | int null | numeric(10) | Indicates what type of privileges to the custom property value will be assigned to a user or group who has Viewer access rights to the Item object in which the custom property appears. |
| INA_AUTHOR_ACCESS (Author Privileges) | int null | numeric(10) | Indicates what type of privileges to the custom property value will be assigned to a user or group who has Author access rights to the Item object in which the custom property appears. |
| INA_OWNER_ACCESS (Owner Privileges) | int null | numeric(10) | Indicates what type of privileges to the custom property value will be assigned to a user or group who has Owner access rights to the Item object in which the custom property appears. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| INA_ADMIN_ACCESS <i>(Admin Privileges)</i> | int null | numeric(10) | Indicates what type of privileges to the custom property value will be assigned to a user or group who has Admin access rights to the Item object in which the custom property appears. |
| INA_ATTR_NO <i>(Target Property Number)</i> | int null | numeric(10) | The number of the property in the object for which this particular custom property is being specified. |
| INA_INPUT_REQUIRED <i>(Input Required?)</i> | int null | numeric(10) | Indicates whether a value must be specified for this particular custom property. |
| INA_DEFAULT_VALUE <i>(Default Value)</i> | varchar(254) null | varchar2(254) null | Provides an initial value (currently 64 bytes) for this particular custom property. |
| INA_INDEX <i>(Indexed?)</i> | tinyint null | numeric(3) null | Contains a flag for indexing this property in the library system. |
| INA_QUICK_LIST <i>(Quick List?)</i> | tinyint null | numeric(3) null | Contains a flag to create an unsecured table (look aside list) of distinct values for this property. This property can be viewed, enabled, or disabled by administrative users. |
| INA_PARENT_OBJECT <i>(Parent Object Number)</i> | int null | numeric(10) | Contains the object number of the parent for which this (child) custom property object is being created. |
| INA_MULTI_VALUED <i>(Multi-valued?)</i> | tinyint null | numeric(3) | Indicates whether this custom property is multi-valued or single-valued. |
| INA_NODISPLAY <i>(Hide?)</i> | tinyint null | numeric(3) | Indicates whether the custom property display name and value are hidden (Yes) from or displayed to (No) users in the associated Item, Version, or Checkout object. The default is No. Setting this value to Yes hides the custom property from all users, including administrators. This property, when set to Yes, provides a way for administrators to turn off the use of a custom property or, more likely, for programs to create custom properties that are not for user input or viewing in the first place. |

ITEM_MV_CUSTOM01 Table

ITEM_MV_CUSTOM01 Table (SPI_OBJ_ITEM_MV_CUSTOM01, 10240)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV01_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV01_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM02 Table

ITEM_MV_CUSTOM02 Table (SPI_OBJ_ITEM_MV_CUSTOM02, 10496)

There are 10 possible Item Multi valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV02_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV02_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM03 Table

ITEM_MV_CUSTOM03 Table (SPI_OBJ_ITEM_MV_CUSTOM03, 10752)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV03_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV03_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM04 Table

ITEM_MV_CUSTOM04 Table (SPI_OBJ_ITEM_MV_CUSTOM04, 11008)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV04_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV04_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM05 Table

ITEM_MV_CUSTOM05 Table (SPI_OBJ_ITEM_MV_CUSTOM05, 11264)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV05_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV05_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM06 Table

ITEM_MV_CUSTOM06 Table (SPI_OBJ_ITEM_MV_CUSTOM06, 11520)

There are 10 possible Item Multi valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV06_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV06_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM07 Table

ITEM_MV_CUSTOM07 Table (SPI_OBJ_ITEM_MV_CUSTOM07, 11776)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV07_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV07_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM08 Table

ITEM_MV_CUSTOM08 Table (SPI_OBJ_ITEM_MV_CUSTOM08, 12032)

There are 10 possible Item Multi valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV08_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV08_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM09 Table

ITEM_MV_CUSTOM09 Table (SPI_OBJ_ITEM_MV_CUSTOM09, 12288)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV09_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV09_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

ITEM_MV_CUSTOM10 Table

ITEM_MV_CUSTOM10 Table (SPI_OBJ_ITEM_MV_CUSTOM10, 12544)

There are 10 possible Item Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ITMV10_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Reference to Item ID property (E_NAME) of Item object instance. |
| ITMV10_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a user supplied value that must meet type and length restrictions. |

LANGID Table

LANGID Table (SPI_OBJ_LANGUAGE, 5376)

(Internal use only) Used for national language support (NLS).

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| LANG_ID (<i>System Name</i>) | tinyint null | numeric(3) null | Contains the display language ID: English 1 French 2 German 3 Korean 24 Japanese 25 |

LIBRARY Table

LIBRARY Table (SPI_OBJ_SYSTEM, 1280)

Describes the library system elements (for example, users, groups, items, custom properties) and supplies system-wide limits and defaults.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| L_NAME (<i>System Name</i>) | varchar(32) | varchar2(32) | Contains the name assigned to the library system when it was installed. |
| L_INST_NAME (<i>Installer</i>) | varchar(32) null | varchar2(32) null | Contains the name of the person who installed the library system. |
| L_INST_DATE (<i>Installation Date</i>) | datetime null | date null | Contains the date and time when the library system was installed. |
| L_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about the Content Services system. |
| L_SW_VERSION (<i>Content Services Software Version</i>) | varchar(32) null | varchar2(32) null | Contains the software version of the product being used (as specified in the L_PRODUCT property in this object). |
| L_VARIANT_FLAG (<i>Default Allow Secondary Lines?</i>) | int null | numeric(10) null | Indicates whether secondary version lines are allowed for items in this system. |
| L_VERSION_LIMIT (<i>Default Online Limit</i>) | int null | numeric(10) null | Provides an initial value for the E_ONLINE_LMT property in the Item object created when a document is added to the library system. |
| L_DISPOSITION (<i>Default Archive Category</i>) | varchar(32) null | varchar2(32) null | Provides a value for the SPI_SES_DEF_DISP_CATEGORY property in the Session object while the session is open and active. |
| L_CHKOUT_F_NAME (<i>Default Checkout File Name</i>) | int null | numeric(10) null | Indicates which option the system is to use to create the initial value for the SPI_ITEM_CHECKOUT_FILE property in a new Item object that is created when a file is being added. |
| L_ENABLE_REPL (<i>Replication Enabled?</i>) | int null | numeric(10) null | Indicates whether replication is enabled for this system. |
| L_DEFAULT_ACCESS (<i>Default General Access</i>) | int null | numeric(10) null | Not Used. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|---------------------|----------------------|---|
| L_STOR_ALLOC (Default Storage Category) | varchar(32) null | varchar2(32) null | Contains the value to apply to the SPI_SES_DEF_FILE_CATEGORY property in the Session object while the session is open and active. |
| L_LOGIN_COUNT (Number of Active Sessions) | int null | numeric(10) null | Contains the current number of active sessions. |
| L_USER_COUNT (Number of Users) | int null | numeric(10) null | Contains the total number of users who have been added to the library system. |
| L_BACKUP_DEVICE (Backup Device) | varchar(32) null | varchar2(32) null | Reserved for future use. |
| L_BACKUP_PATH (Backup Path) | varchar(96) null | varchar2(96) null | Reserved for future use. |
| L_SERIAL_NO (Serial Number) | varchar(32) null | varchar2(32) null | Contains the serial number of the license for the library system. |
| L_NETWORK_LOGIN (Obsolete) | Obsolete | Obsolete | Obsolete |
| L_INST_TYPE (System Type) | int null | numeric(10) null | Indicates the library system type. |
| L_CMD_LINE_PSWD (Allow Cmd-line Passwd?) | int null | numeric(10) null | No longer used. |
| L_LOGINS_ENABLED (Logons Enabled?) | int null | numeric(10) null | Indicates whether users whose default group is not set to Administrators can open sessions. |
| L_PRIMARY_SERVER (Primary Storage Manager) | varchar(32) null | varchar2(32) null | Provides a value for the SE_PRIMARY_SERVER property in the Session object while the session is open and active. The value indicates the Storage Manager that Content Services will attempt to use first to handle an operational request from a user. |
| L_THING_1 (Future) | varchar(32) null | varchar2(32) null | For future use. |
| L_THING_2 (Future) | varchar(32) null | varchar2(32) null | For future use. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|-------------------|---------------------|--|
| L_PV_IN_LIMIT <i>(Permanent Versions in Limit?)</i> | int null | numeric(10) null | Indicates whether versions marked as Permanent are to be included in the online limit for a protected item. |
| L_SECURE_SEARCH <i>(Secure Searches?)</i> | int null | numeric(10) null | Indicates whether users must have at least Viewer access rights to an item to have that item included in the results of their searches. |
| L_SUSPENDED_COUNT <i>(Num. of Suspended Sess)</i> | int null | numeric(10) null | Contains the current number of sessions that are suspended. |
| L_CS_ENABLED <i>(Content Search Enabled?)</i> | int null | numeric(10) null | Indicates whether content searches are allowed. |
| L_CSINDEX_DEFAULT <i>(Default Index Versions?)</i> | tinyint null | numeric(3) null | Indicates a default value for the Default Index Versions? property in Item objects. |
| L_CSINDEX_LIMIT <i>(Default Indexed Version Limit)</i> | smallint null | numeric(5) null | Indicates a default value for the Indexed Version Limit property in Item objects |
| L_CSI_PV_IN_LIMIT <i>(Permanent Indexes in Limit?)</i> | tinyint null | numeric(3) null | Indicates whether versions marked as permanently indexed (in the SPI_VERS_CSI_PERMANENT property) are to be included in the indexed version limit for an item. |
| L_SEARCH_DOMAIN <i>(Access Domain)</i> | int null | numeric(10) null | <p>Contains one of the following group-domain values indicating the group affiliation(s) to apply when determining which search results to return to the caller during a search:</p> <p>SPI_ACCDOMAIN_ACTIVEGROUP_S</p> <p>Search results based on the access rights assigned to the user active group (the default value).</p> <p>SPI_ACCDOMAIN_ALLGROUPS_S</p> <p>Search results based on the highest access rights assigned to any of the groups in which the user is a member.</p> |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|---------------------|----------------------|---|
| L_ADMIN_USER (Administrative User) | varchar(32) null | varchar2(32) null | The name given to the initial administrator in the particular Content Services system, with the default value being Admin. This user is added automatically by the installation program. |
| L_ADMIN_GROUP (Administrative Group) | varchar(32) null | varchar2(32) null | The name given to the administrative group, with the default value being Administrators. |
| L_GENERAL_GROUP (Universal Group) | varchar(32) null | varchar2(32) null | The name given to the universal group, with the default value being General Users. |
| L_LANGUAGE (Display Language) | tinyint null | numeric(3) null | Contains the language in which the software is being presented (for example, German). |
| L_ROOT (Root Path) | varchar(32) null | varchar2(32) null | Contains the path name of the root directory used for the Content Services installation. |
| L_PRODUCT (Product Name) | varchar(32) null | varchar2(32) null | Contains the name to be used for the ELS product. The default value is <i>IDMDS</i> . |
| L_SEQ_ITEM_IDS (Sequential Item Ids?) | tinyint null | numeric(3) null | Specifies whether to use sequential item ID rather than Julian date ID for all items added. |
| L_CS_THRESHOLD (Content Search Threshold) | int null | numeric(10) null | Determines whether to do a content search or property search first, when both types of criteria are specified for a search. |
| L_RECLAIM_ACCESS (Min Reclaim Access Level) | smallint null | numeric(5) null | Indicates the minimum level of access rights needed to reclaim an item. |
| L_ITEM_DEL_ACCESS (Min Item Delete Access Level) | smallint null | numeric(5) null | Indicates the minimum level of access rights needed to delete an item. The only acceptable values are Admin or Owner access. NOTE: Regardless of this setting, if an item has no associated versions, anyone with Owner access rights to the Item object can delete it. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| L_WORKFLOW_ON <i>(Workflow on?)</i> | int null | numeric(10) null | Indicates if Workflow is on. |
| L_CUSTODIAL_SERVER <i>(Database Custodial Server)</i> | varchar(32) null | varchar2(32) null | Provides a value for the L_CUSTODIAL_SERVER property in the Session object while the session is open and active. The value is the name of the storage server that contains the Storage Manager in charge of running the database usage and session cleanup daemons. When changed to a different storage server name, the Storage Manager on the newly assigned server will pick up the task on the next daemon cycle. |
| L_CSI_THESAURUS <i>(Thesaurus in Content Search?)</i> | tinyint(1) | tinyint(1) | Indicates whether CBR searches will include a thesaurus or not. By default, no thesaurus is included. |
| L_VERITY_ADMIN <i>(Verity Admin Server)</i> | varchar(32) null | varchar(32) null | The server designated as the K2 Master Admin for the library (set during installation). |
| L_VERITY_ADMINPORT <i>(Verity Admin Server Port)</i> | varchar(32) null | varchar(32) null | The port assigned to the K2 Master Admin (set during installation). |
| L_NETLOGIN_CI <i>(Case Insensitive Network Login?)</i> | varchar(32) null | varchar(32) null | Case-insensitive network logon name |

MESSAGE Table

MESSAGE Table (SPI_OBJ_MESSAGE, 6656)

Contains a message that is typically used to provide information, warn of pending problems, or identify error conditions. Message objects can be child objects of User or Group objects.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| M_TO (To) | varchar(32) null | varchar2(32) null | Contains the name of the user or group that is the intended receiver of the message. |
| M_OBJ_TYPE (Parent Object Type) | int null | numeric(10) null | Contains the name of the parent User or Group object where the message object is to be listed. At least Viewer access to the parent object instance is required. |
| M_TYPE (Type) | int null | numeric(10) null | Contains the type of message (where: SPI_MESSAGE_INFO_S = Informational; SPI_MESSAGE_WARNING_S = Warning; SPI_MESSAGE_ERROR = Error). |
| M_FROM (From) | varchar(32) null | varchar2(32) null | Indicates the source of the message. |
| M_DATE (Date) | datetime null | date null | Contains the date and time that the message was sent. |
| M_TEXT (Message) | varchar(224) null | varchar2(224) null | Contains the text of the message. |
| M_ID_NUM (Message ID) | int null | numeric(10) | Contains the unique numerical identifier of the message. |

NAME_SERVICE Table

NAME_SERVICE Table (SPI_OBJ_NAME_SERVICE, 4864)

Provides client software with connection information needed to connect to a Storage or Content Search Manager.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|-----------------------|------------------------|--|
| NS_NODE_NAME <i>(Node Name)</i> | varchar (32) | varchar2 (32) | Specifies node name where service is running. |
| NS_SERVICE_TYPE <i>(Service Type)</i> | varchar (32) | varchar2 (32) | Specifies the service type (SHF / CSMGR). |
| NS_PROTOCOL <i>(Protocol)</i> | varchar (32) | varchar2 (32) | Specifies the protocol used by the service. |
| NS_CONNECTION_ADDR <i>(Connection Address)</i> | varchar (290) null | varchar2 (290) null | Specifies the connection information, including port, for the service. |
| NS_COST <i>(Connection Cost)</i> | int null | numeric (10) null | Specifies the cost of service. |
| NS_ENDIAN | int null | numeric (10) null | Specifies endianness of the service. |
| NS_CHARSET <i>(Character Set)</i> | varchar (128) null | varchar2 (128) null | Default character set for the service. |

NUMID Table

NUMID Table

Lists the last number used for sequentially numbered objects such as relationships, sessions.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|-------------------|----------------------|-----------------------|--|
| ST_ID_NUM | int null | numeric (10) null | Contains last number used for ST_ID attribute in STACK table. |
| SE_ID_NUM | int null | numeric (10) null | Contains last number used for SE_ID attribute in SESSION table. |
| DS_ID_NUM | int null | numeric (10) null | Contains last number used for DS_ID attribute in DISP_STACK table. |
| M_ID_NUM | int null | numeric (10) null | Contains last number used for M_ID_NUM attribute in MESSAGE table. |
| CS_ID_NUM | int null | numeric (10) null | Contains last number used for CSI_ID attribute in CSI_INDEX table. |
| REL_ID_NUM | int null | numeric (10) null | Contains last number used for REL_ID attribute in RELATION table. |
| WL_ID_NUM | int null | numeric (10) null | Contains last number used for WL_ID attribute in WORKFLOW_LOG table. |

OBJECT_INFO Table

OBJECT_INFO Table (SPI_OBJ_OBJ_INFO, 18688)

Provide the list of the objects in Content Services property repository.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---------------------------------|-----------------------|------------------------|------------------------|
| OI_ID (Object ID) | int | numeric (10) | Specifies Object ID |
| OI_NAME (Object Name) | varchar (128) null | varchar2 (128) null | Specifies Object Name. |

OTHERINST Table

OTHERINST Table (SPI_OBJ_OTHERSYS, 8704)

Identifies another system of Content Services in which users can open sessions from their current system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| OI_NAME (<i>System Name</i>) | varchar(32) | varchar2(32) | Contains the name assigned to this other system when it was installed. |
| OI_NODE (<i>Property Server</i>) | varchar(32) null | varchar2(32) null | Contains the name of the server on which this other system is installed. |
| OI_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about this other system. |

RELATION Table

RELATION Table (SPI_OBJ_RELATIONSHIP, 9984)

Describes a particular relation between objects (for example, between parent and child Version objects in a compound document). Note that this object is primarily for FileNet internal use, but it is a free-form object available for customization of user interfaces built with the CS Toolkit. Therefore, it is typically not useful for creating reports unless you are using a customized user interface and know specifically how that interface utilizes the Relation object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| REL_ID <i>(Relationship ID)</i> | int null | numeric(10) null | Unique Relationship object identifier. |
| REL_PARENT_OBJ <i>(Parent Object Type)</i> | int null | numeric(10) null | The object type that is the parent for the particular relationship. |
| REL_PARENT <i>(Parent Object Key)</i> | varchar(82) | varchar2(82) | The key of the Parent object instance for the particular relationship, as returned by spiKeyBuild. |
| REL_CHILD_OBJ <i>(Child Object Type)</i> | int null | numeric(10) null | The object type that is the parent for the particular relationship. |
| REL_CHILD <i>(Child Object Key)</i> | varchar(82) | varchar2(82) | The key of the Child object instance for the particular relationship, as returned by spiKeyBuild. |
| REL_OWNER_OBJ <i>(Owner Object Type)</i> | int null | numeric(10) null | The object type that owns the particular relationship, which can be used for extending relationship security model. |
| REL_OWNER <i>(Owner Object Key)</i> | varchar(82) | varchar2(82) | The key of the object instance that owns the particular relationship, as returned by spiKeyBuild. |
| REL_CLASS <i>(Relationship Class)</i> | varchar(32) null | varchar2(32) null | The class of the particular relationship. |
| REL_CLASS_PROPS <i>(Relationship Class Properties)</i> | varchar(32) null | varchar2(32) null | The property values of the relationship class identified in SPI_REL_CLASS. You can use these properties to further qualify the particular relationship class. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| REL_SEQNO <i>(Relationship Sequence ID)</i> | int null | numeric(10) null | The sequence ID of the particular relationship. This property can be used to present a series of relationship in a particular sequential order; however, there is no behavior enforced on this property. |
| REL_STRENGTH <i>(Relationship Strength)</i> | int null | numeric(10) null | The strength of the particular relationship. This can be used to qualify relationships, such as those required for compound documents. |
| REL_DYN_PROP_ID <i>(Dynamic Property ID)</i> | int null | numeric(10) null | For internal FileNet use. |
| REL_DYN_PROP_VALUE <i>(Dynamic Property Value)</i> | varchar(64) null | varchar2(64) null | For internal FileNet use. |
| REL_REPL_STATUS <i>(Relationship Replication Info)</i> | int null | numeric(10) null | For internal FileNet use. |
| REL_REL_INTEGER1-10 <i>(Integer 1-10)</i> | int null | numeric(10) null | Contains one to ten customized integer properties (each up to 9 characters) that describe the object. |
| REL_REL_STRING1-5 <i>(String 1-5)</i> | varchar(32) null | varchar2(32) null | Contains one to five customized string properties (each up to 32 characters) that describe the object. |
| REL_REL_MSTRING1-5 <i>(Medium String 1-5)</i> | varchar(64) null | varchar2(64) null | Contains one to five customized medium-length string properties (each up to 64 characters) that describe the object. |
| REL_REL_LSTRING1-3 <i>(Long String 1-5)</i> | varchar(254) null | varchar2(254) null | Contains a comment about the archive repository. |

REPL_KW Table

REPL_KW Table (SPI_OBJ_KEYWORD, 4096)

Provides information about a keyword that is associated with a particular item. Keyword objects appear as child objects of an Item object. Stores the information about the unique replicated folder tree in which the item is filed.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--------------------------------------|----------------------|-----------------------|---|
| KW_E_NAME <i>(Item ID)</i> | varchar (9) | char (9) | Specifies Item ID |
| KW_NAME <i>(Folder ID)</i> | varchar (244) | varchar2 (244) | Specifies the folder ID of the replicated tree where item is filed. |

REPLOBJ Table

REPLOBJ Table (SPI_OBJ_REPL_DATA, 17920)

Used to create queue of replication work to be performed by remote servers.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|-----------------------|------------------------|---|
| RD_WORK_ID (<i>Replication request Work ID</i>) | varchar (254) | varchar2 (254) | Replication Work ID |
| RD_TYPE (<i>Replication Work Type</i>) | smallint null | numeric (5) null | Specifies type of replication request. |
| RD_STATUS (<i>Replication Status</i>) | smallint null | numeric (5) null | Specifies the status of the replication request. |
| RD_DATETIME (<i>Work request date</i>) | datetime null | date null | Specifies date and time of replication request. |
| RD_PARENT (<i>Parent Folder ID</i>) | varchar (254) null | varchar2 (254) null | Specifies parent folder ID. |
| RD_OBJECT (<i>replicating object</i>) | varchar (254) null | varchar2 (254) null | Specifies replicating object. |
| RD_TID (<i>Tree ID</i>) | varchar (254) null | varchar2 (254) null | Specifies ID of the root replica folder of the item. |
| RD_DATA (<i>Raw replication data</i>) | varchar (254) null | varchar2 (254) null | Specifies item reference count, replication work type, parent folder ID and timestamp values. |
| RD_SRC_REVISION | int null | numeric (10) null | Source revision of an object. |
| RD_TARG_REVISION | int null | numeric (10) null | Target revision of an object. |
| RD_STRING1 | varchar (32) null | varchar2 (32) null | Contains version information |
| RD_LSTRING1 | varchar (254) null | varchar2 (254) null | Contains Home Server status information. |
| RD_INTEGER1 | int null | numeric (10) null | |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| RD_ERROR_COUNT (Replication Error Count) | int null | numeric (10) null | Contains the number of errors occurred during processing of replication requests. |
| RD_LAST_ERROR (<i>Last Error Number</i>) | int null | numeric (10) null | Contains the last error encountered while processing replication requests. |

SEARCH_REQUEST Table

SEARCH_REQUEST Table (SPI_OBJ_SEARCH_REQUEST, 1024)

Describes a search, which can return a set of objects that match a simple or complex query description.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|------------------------|---|
| SR_NAME (<i>Search Name</i>) | varchar(32) null | varchar2(32) null | Contains the unique name of the search. |
| SR_TYPE (<i>Search Type</i>) | int null | numeric(10) null | Indicates the type of parent object (SPI_OBJ_USER, SPI_OBJ_GROUP, or SPI_OBJ_SYSTEM) named in the SR_REF property. |
| SR_REF (<i>Search Reference</i>) | varchar(32) null | varchar2(32) null | Contains the name of the parent User, Group, or System object for the Search object. |
| SR_OWNER (<i>Added By User</i>) | varchar(32) null | varchar2(32) null | Contains the user name of the person who added the Search object to the property repository. |
| SR_CREATE_DATE (<i>Date Added</i>) | datetime null | date null | Contains the date and time that the Search object was added to the property repository. |
| SR_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about the search. |
| SR_QUERY (<i>Query</i>) | text null | long null | Contains the query statements that describe the results to be returned by the search. |
| SR_SHOW (<i>Show</i>) | varchar(254) null | varchar2(254) null | Contains the names of the properties that will appear in the partial OAB that result from the Search query described in the SR_QUERY property. |
| SR_SORT (<i>Sort</i>) | varchar(254) null | varchar2(254) null | Contains the sort order (SPI_QUERY_ASCENDING or SPI_QUERY_DESCENDING) for the partial OAB that result from the Search query described in the SR_QUERY property. |
| SR_QUERY_CS (<i>Content Search Query</i>) | text null | varchar2(1092) null | Contains a long string (1024 bytes) containing the content query specification. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| SR_SEARCH_DOMAIN <i>(Access Domain)</i> | int null | numeric(10) null | <p>Contains one of the following group-domain values indicating the group affiliation(s) to apply when determining which search results to return to the caller during a search:</p> <p>SPI_ACCDOMAIN_DEFAULT_S</p> <p>Search results based on the first non NULL access domain value passed from the user Session, User, active Group, or System object (assessed in that order).</p> <p>SPI_ACCDOMAIN_ACTIVEGROUP_S</p> <p>Search results based on the active user group.</p> <p>SPI_ACCDOMAIN_ALLGROUPS_S</p> <p>Search results based on all groups in which the user is a member.</p> |
| SR_ACCESS_LEVEL <i>(Access Level)</i> | smallint null | numeric(5) null | <p>Indicates the access level to apply when determining which search results to return to the caller during a search (assuming SPI_SYS_SECURE_SEARCH is set to SPI_YESNO_YES_S in the System object).</p> |
| SR_COUNT <i>(Count Only)</i> | tinyint null | numeric(3) null | <p>Indicates whether to return the number of search results, rather than the actual results themselves, during a search.</p> |
| SR_DISTINCT <i>(Distinct?)</i> | tinyint null | numeric(3) null | <p>Indicates whether to return only the distinct, unduplicated rows, rather than all rows meeting the search criteria.</p> |
| SR_SECURE_SEARCH <i>(Secure Searches?)</i> | int null | numeric(10) null | <p>Indicates whether users must have at least Viewer access rights to an item to have that item included in the results of their searches.</p> |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| SR_CS_THRESHOLD <i>(Content Search Threshold)</i> | int null | numeric(10) null | Determines whether to do a content search or property search first, when both types of criteria are specified for a search. |

SESSIONTAB Table

SESSIONTAB Table (SPI_OBJ_SESSION, 2816)

Describes a open user session of Content Services. The system adds a Session object each time a Content Services session is opened.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| SE_ID (<i>Session Number</i>) | int | numeric(10) | Contains the unique identifier for the open session. |
| SE_OWNER (<i>User Name</i>) | varchar(32) null | varchar2(32) null | Contains the user name of the person who opened the session. |
| SE_GROUP (<i>Active Group</i>) | varchar(32) null | varchar2(32) null | Contains the name of the active user group. |
| SE_DATE (<i>Start Date</i>) | datetime null | date null | Contains the date and time the user opened the session. |
| SE_WS_ADR (<i>Client Workstation Address</i>) | varchar(32) null | varchar2(32) null | Contains the network address of the workstation on which the user opened the session. |
| SE_ADMIN (<i>Admin Privileges?</i>) | int null | numeric(10) null | (Internal use only) Indicates whether the user was in the Administrators group when he or she opened the session. |
| SE_FILE_C (<i>Add Items?</i>) | int null | numeric(10) null | Indicates whether the user who opened the session can add and annotate Item objects in the property repository. |
| SE_USER_C (<i>Add Users?</i>) | int null | numeric(10) null | Indicates whether the user who opened the session can add User objects to the property repository. |
| SE_GRP_C (<i>Add Groups?</i>) | int null | numeric(10) null | Indicates whether the user who opened the session can add Group objects to the property repository. |
| SE_PATH (<i>Default Checkout Path</i>) | varchar(96) null | varchar2(96) null | Contains the value from the U_PATH property in the User object of the user who opened the session. |
| SE_F_EXT (<i>Obsolete</i>) | Obsolete | Obsolete | Obsolete |
| SE_STOR_ALLOC (<i>Default Storage Category</i>) | varchar(32) null | varchar2(32) null | Contains the default storage category that will be assigned to any new items added by the user during this session. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|---------------------|----------------------|---|
| SE_DISPOSITION <i>(Default Archive Category)</i> | varchar(32) null | varchar2(32) null | Contains the default archive category that will be assigned to any new items added by the user during this session. |
| SE_STATE <i>(Status)</i> | int null | numeric(10) null | Indicates the status of the session. |
| SE_SUSPEND_DATE <i>(Suspension Date)</i> | datetime null | date null | Indicates the date and time the session was suspended. |
| SE_CATALOG_ID <i>(Property Manager Process ID)</i> | int null | numeric(10) null | Contains the unique identifier of the Property Manager process that is servicing the session. |
| SE_PRIMARY_SERVER <i>(Primary Storage Manager)</i> | varchar(32) null | varchar2(32) null | Contains the primary Storage Manager for operational requests submitted by the user during this session. |
| SE_INSTALLATION <i>(System Name)</i> | varchar(32) null | varchar2(32) null | Contains the name of the library system against which this session was opened. |
| SE_SECURE_SEARCH <i>(Secure Searches?)</i> | int null | numeric(10) null | Indicates whether the user who opened this session must have at least Viewer access rights to an item to have that item included in the results of one of his or her searches. |
| SE_SEARCH_DOMAIN <i>(Access Domain)</i> | int null | numeric(10) null | <p>Contains one of the following group domain values indicating the group affiliation(s) to apply when determining which search results to return:</p> <p>SPI_ACCDOMAIN_ACTIVEGROUP_S</p> <p>Search results based on the user active group.</p> <p>SPI_ACCDOMAIN_ALLGROUPS_S</p> <p>Search results based on all groups in which the user is a member.</p> |
| SE_CSINDEX_DEFAULT <i>(Default Index Versions?)</i> | tinyint null | numeric(3) null | Contains the value that will be entered for the Default Index Versions? property in new Item objects added by the user during this session. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| SE_LANGUAGE <i>(Display Language)</i> | tinyint null | numeric(3) null | Contains the language with which the user opened the session. |
| SE_NET_USER <i>(Network Login Name)</i> | varchar(224) null | varchar2(224) null | Contains the network login name used to open the session. |
| SE_CS_THRESHOLD <i>(Content Search Threshold)</i> | int null | numeric(10) null | Determines whether to do a content search or property search first, when both types of criteria are specified for a search. |
| SE_RECLAIM_ACCESS <i>(Min Reclaim Access Level)</i> | smallint null | numeric(5) null | Indicates the minimum level of access rights needed to reclaim an item. |
| SE_ITEM_DEL_ACCESS <i>(Min Item Delete Access Level)</i> | smallint null | numeric(5) null | Indicates the minimum level of access rights needed to delete an item. The value of this property is supplied from the L_ITEM_DEL_ACCESS property of the System object. This property value is supplied when the session opens. Any changes to L_ITEM_DEL_ACCESS are not effected until the next session is opened. |

STACK Table

STACK Table (SPI_OBJ_SHELF, 2304)

Describes a storage repository, which is used to store version files managed by Content Services.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ST_ID (<i>Storage Repository ID</i>) | int | numeric(10) | Contains the unique numeric identifier supplied by Content Services for the storage repository. |
| ST_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about the storage repository. |
| ST_SPACE (<i>Max Disk Space</i>) | int null | numeric(10) null | Indicates the limit on disk space used by files on this storage repository. |
| ST_USED_SPACE (<i>Total Space Used</i>) | int null | numeric(10) null | Indicates the current disk space used by this storage repository. |
| ST_FILES (<i>Total Files</i>) | int null | numeric(10) null | Indicates the current count of files on this storage repository. |
| ST_NODE (<i>Server</i>) | varchar(32) null | varchar2(32) null | Contains the unique network name identifying the server on which the Storage Manager identified by the ST_SERVR_NM property is located. |
| ST_DEV (<i>Device</i>) | varchar(32) null | varchar2(32) null | Contains the name of the storage device where the storage repository is located. |
| ST_DIR (<i>Directory</i>) | varchar(32) null | varchar2(32) null | Contains the immediate directory name used to locate the storage repository relative to the node, device, and path. |
| ST_SERVR_NM (<i>Server Name</i>) | varchar(32) null | varchar2(32) null | Contains the Content Services process name for the storage server on which this storage repository is located. |
| ST_SERVR_CLASS (<i>Server Class</i>) | varchar(32) null | varchar2(32) null | (Internal use only) Contains the class (for example, Windows) of the server on which this storage repository is located. |
| ST_PATH (<i>Path</i>) | varchar(96) null | varchar2(96) null | Contains the path location of the storage repository. |
| ST_THING_1 (<i>Future</i>) | varchar(32) null | varchar2(32) null | Reserved for future use. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| ST_THING_2 <i>(Future)</i> | varchar(32) null | varchar2(32) null | Reserved for future use. |
| ST_NET_PATH <i>(Network Path)</i> | varchar(224) null | varchar2(224) null | Contains the network path location of a storage repository located on a network server. |
| ST_MAX_FILES <i>(Max Files)</i> | int null | numeric(10) null | Indicates the maximum number of files allowed on the storage repository. |
| ST_WARN_FILES <i>(Max Files Warning Threshold)</i> | int null | numeric(10) null | (Currently unimplemented) Indicates the warning threshold for the file count limit. |
| ST_WARN_SPACE <i>(Max Space Warning Threshold)</i> | int null | numeric(10) null | Currently unimplemented) Indicates the warning threshold for disk space limit. |
| ST_OVERFLOW_ID <i>(Overflow Storage Repository)</i> | int null | numeric(10) null | (Currently unimplemented) Indicates the storage repository to which new files will be directed when a limit is reached. |

STORAGE_ALLOC Table

STORAGE_ALLOC Table (SPI_OBJ_FILE_CATEGORY, 2048)

Describes a storage category, which is used to categorize particular Items and to assign their version files to a common storage repository.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| SA_ALLOC_NAME <i>(Storage Category)</i> | varchar(32) | varchar2(32) | Contains the name of the storage category. |
| SA_ST_ID <i>(Assigned Storage Repository)</i> | int null | numeric(10) null | Contains the unique numeric identifier of the storage repository to which the storage category is linked. |
| SA_COMMENT <i>(Comment)</i> | varchar(224) null | varchar2(224) null | Contains a comment about the storage category. |
| SA_CSI_ID <i>(Assigned Content Search Repository)</i> | int null | numeric(10) null | Contains the ID of the content search repository assigned to this storage category. |

USR Table

USR Table (SPI_OBJ_USER, 1792)

Describes a user who has been added to the library system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| U_NAME <i>(User Name)</i> | varchar(32) null | varchar2(32) null | Contains the name of the user. |
| U_CREATE_DATE <i>(Date Added)</i> | datetime null | date null | Contains the date and time the user was added to the library system. |
| U_COMMENT <i>(Comment)</i> | varchar(224) null | varchar2(224) null | Contains a comment about the user. |
| U_PWD <i>(Password)</i> | varchar(32) null | varchar2(32) null | Contains an optional password for the user. If any value is present, a password must be provided to open a session. If the value is blank (contains only a carriage return) then no password is required to open a session. |
| U_ACT_GRP <i>(Default Group)</i> | varchar(32) null | varchar2(32) null | Contains the name of the group to be used as the active group when the user opens a session. |
| U_PATH <i>(Default Path)</i> | varchar(96) null | varchar2(96) null | Provides a value for the SE_PATH in the Session object. This value contains the default path (drive and directory names only) that indicates where version files will be sent when they are checked out or copied by this user. |
| U_FILE_CREATOR <i>(Add Items?)</i> | int null | numeric(10) null | Provides the value for the SPI_SES_ADD_ITEMS property in the Session object, which indicates whether the user can add and annotate items in the library system. |
| U_USER_CREATOR <i>(Add Users?)</i> | int null | numeric(10) null | Provides the value for the SPI_SES_ADD_USERS property in the Session object, which indicates whether the user can add users to the library system. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|---------------------|----------------------|---|
| U_GROUP_CREATOR <i>(Add Groups?)</i> | int null | numeric(10) null | Provides the value for the SPI_SES_ADD_GROUPS property in the Session object, which indicates whether the user can add groups to the library system. |
| U_STOR_ALLOC <i>(Default Storage Category)</i> | varchar(32) null | varchar2(32) null | Contains the value to apply to the SPI_SES_DEF_FILE_CATEGORY property in the Session object while the session is open and active. |
| U_DISPOSITION <i>(Default Archive Category)</i> | varchar(32) null | varchar2(32) null | Provides a value for the SPI_SES_DEF_DISP_CATEGORY property in the Session object while the session is open and active. |
| U_CHKOUT_F_name <i>(Obsolete)</i> | Obsolete | Obsolete | Obsolete |
| U_CHKOUT_F_EXT <i>(Obsolete)</i> | Obsolete | Obsolete | Obsolete |
| U_OWNER_NAME <i>(Added By User)</i> | varchar(32) null | varchar2(32) null | Contains the user name of the person who added the User object to the library system. |
| U_NET_NAME <i>(Network Logon Name)</i> | varchar(64) null | varchar2(64) null | Contains the network login name for the user. |
| U_PRIMARY_SERVER <i>(Primary Storage Manager)</i> | varchar(32) null | varchar2(32) null | Provides a value for the SE_PRIMARY_SERVER property in the Session object while the session is open and active. The value indicates the Storage Manager that Content Services will attempt to use first to handle an operational request from the user. |
| U_STATUS <i>(Enabled?)</i> | int null | numeric(10) null | Indicates whether a user may open a Content Services session. |
| U_DEF_RUN_APP <i>(Default Run Application?)</i> | int null | numeric(10) null | Indicates whether the application associated with a version will be run by default when that document is checked out or copied. |
| U_CSINDEX_DEFAULT <i>(Default Index Versions?)</i> | tinyint null | numeric(3) null | Indicates a default value for the Default Index Versions? property in Item objects. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|---|
| U_SEARCH_DOMAIN <i>(Access Domain)</i> | int null | numeric(10) null | <p>Contains one of the following group domain values indicating the group affiliation(s) to apply when determining which search results to return during the search:</p> <p>SPI_ACCDOMAIN_DEFAULT_S</p> <p>Search results based on the first Active Group or All Groups access domain value passed from the user active Group or System object (assessed in that order).</p> <p>SPI_ACCDOMAIN_ACTIVEGROUP_S</p> <p>Search results based on the access rights assigned to the user active group.</p> <p>SPI_ACCDOMAIN_ALLGROUPS_S</p> <p>Search results based on the highest access rights assigned to any groups in which the user is a member.</p> |
| U_LANGUAGE <i>(Display Language)</i> | tinyint null | numeric(3) null | Indicates the default display language. |
| U_SECURE_SEARCH <i>(Secure Searches?)</i> | tinyint null | numeric(3) null | Indicates whether users must have at least Viewer access rights to an item to have that item included in the results of their searches. |
| U_EMAIL_ADDRESS <i>(Email Address)</i> | varchar(64) null | varchar2(64) null | Reserved for future use. |
| U_FULL_NAME <i>(Full Name)</i> | varchar(224) null | varchar2(224) null | The full name of the user (versus the name under which the user logs on to the CS library system). |
| U_CS_THRESHOLD <i>(Content Search Threshold)</i> | int null | numeric(10) null | Determines whether to do a content search or property search first, when both types of criteria are specified for a search. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| U_SYNC <i>(Imported User Type)</i> | int null | numeric(10) null | Enables synchronization of CS library user names with Windows NT domain(s) user names. The User Import tool (dirsync.exe) runs on the initial storage server as a daemon under a selected library system's Storage Manager. When enabled, the daemon performs ongoing one way synchronization of group and user information from selected network domains to a selected library system database; it does not propagate changes from the library system database back to the network domains. |

VCKO_MV_CUSTOM01 Table

VCKO_MV_CUSTOM01 Table (SPI_OBJ_VCKO_MV_CUSTOM01, 15360)

There are 10 possible Checkout Multiple values for Custom Property objects that can be defined. These objects are created by adding multiple custom property values to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VCKOMV01_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV01_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV01_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV01_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value, supplied by the user that must meet type and length restrictions. |

VCKO_MV_CUSTOM02 Table

VCKO_MV_CUSTOM02 Table (SPI_OBJ_VCKO_MV_CUSTOM02, 15616)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV02_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV02_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV02_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV02_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM03 Table

VCKO_MV_CUSTOM03 Table (SPI_OBJ_VCKO_MV_CUSTOM03, 15872)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV03_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV03_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV03_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV03_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM04 Table

VCKO_MV_CUSTOM04 Table (SPI_OBJ_VCKO_MV_CUSTOM04, 16128)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV04_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV04_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV04_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV04_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM05 Table

VCKO_MV_CUSTOM05 Table (SPI_OBJ_VCKO_MV_CUSTOM05, 16384)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV05_VC_E_NAME (<i>Item ID</i>) | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV05_VC_NAME (<i>Version ID</i>) | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV05_VC_USER (<i>Status</i>) | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV05_VALUE (<i>Custom Property Value</i>) | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM06 Table

VCKO_MV_CUSTOM06 Table (SPI_OBJ_VCKO_MV_CUSTOM06, 16640)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV06_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV06_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV06_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV06_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM07 Table

VCKO_MV_CUSTOM07 Table (SPI_OBJ_VCKO_MV_CUSTOM07, 16896)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV07_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV07_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV07_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV07_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM08 Table

VCKO_MV_CUSTOM08 Table (SPI_OBJ_VCKO_MV_CUSTOM08, 17152)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV08_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV08_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV08_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV08_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM09 Table

VCKO_MV_CUSTOM09 Table (SPI_OBJ_VCKO_MV_CUSTOM09, 17408)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV09_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV09_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV09_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV09_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VCKO_MV_CUSTOM10 Table

VCKO_MV_CUSTOM10 Table (SPI_OBJ_VCKO_MV_CUSTOM10, 17664)

There are 10 possible Checkout Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| VCKOMV10_VC_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (VC_E_NAME) of Checkout object instance. |
| VCKOMV10_VC_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (VC_NAME) of Checkout object instance. |
| VCKOMV10_VC_USER <i>(Status)</i> | varchar(32) null | varchar2(32) | Relates to Checked Out By User property (VC_USER) of Checkout object instance. |
| VCKOMV10_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_DISP Table

VERS_DISP Table (SPI_OBJ_DISP_CATEGORY, 2560)

Describes an archive category that has been added to the library system to identify a method of version archival and to point to an assigned archive repository where the archived versions can be temporarily stored.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|----------------------|-----------------------|--|
| VDO_NAME (<i>Archive Category</i>) | varchar(32) | varchar2(32) | Contains the name of the archive category. |
| VDO_RECLAIM (<i>Reclaimable?</i>) | int null | numeric(10) null | Indicates whether the version can be reclaimed once its file has been moved out of the archive repository. |
| VDO_ST_ID (<i>Assigned Archive Repository</i>) | int null | numeric(10) null | Indicates the name of the archive repository to which the archive category is assigned. |
| VDO_COMMENT (<i>Comment</i>) | varchar(224) null | varchar2(224) null | Contains a comment about the archive category. |

VERS_MV_CUSTOM01 Table

VERS_MV_CUSTOM01 Table (SPI_OBJ_VERS_MV_CUSTOM01, 12800)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV01_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV01_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV01_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM02 Table

VERS_MV_CUSTOM02 Table (SPI_OBJ_VERS_MV_CUSTOM02, 13056)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV02_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV02_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV02_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM03 Table

VERS_MV_CUSTOM03 Table (SPI_OBJ_VERS_MV_CUSTOM03, 13312)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV03_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV03_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV03_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM04 Table

VERS_MV_CUSTOM04 Table (SPI_OBJ_VERS_MV_CUSTOM04, 13568)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV04_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV04_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV04_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM05 Table

VERS_MV_CUSTOM05 Table (SPI_OBJ_VERS_MV_CUSTOM05, 13824)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV05_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV05_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV05_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM06 Table

VERS_MV_CUSTOM06 Table (SPI_OBJ_VERS_MV_CUSTOM06, 14080)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV06_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV06_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV06_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM07 Table

VERS_MV_CUSTOM07 Table (SPI_OBJ_VERS_MV_CUSTOM07, 14336)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV07_V_E_NAME (<i>Item ID</i>) | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV07_V_NAME (<i>Version ID</i>) | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV07_VALUE (<i>Custom Property Value</i>) | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM08 Table

VERS_MV_CUSTOM08 Table (SPI_OBJ_VERS_MV_CUSTOM08, 14592)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV08_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV08_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV08_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM09 Table

VERS_MV_CUSTOM09 Table (SPI_OBJ_VERS_MV_CUSTOM09, 14848)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV09_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV09_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV09_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERS_MV_CUSTOM10 Table

VERS_MV_CUSTOM10 Table (SPI_OBJ_VERS_MV_CUSTOM10, 15104)

There are 10 possible Version Multi-valued Custom Property objects that can be defined. These objects are created by adding a multi-valued custom property to the Custom Property object.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| VERSMV10_V_E_NAME <i>(Item ID)</i> | varchar(9) | char(9) | Relates to Item ID property (V_E_NAME) of the Version object instance. |
| VERSMV10_V_NAME <i>(Version ID)</i> | varchar(7) null | char(7) | Relates to the Version ID property (V_NAME) of the Version object instance. |
| VERSMV10_VALUE <i>(Custom Property Value)</i> | varchar(32) null | varchar2(32) | Contains a value that must meet type and length restrictions. |

VERSION Table

VERSION Table (SPI_OBJ_VERSION, 512)

Describes a version of an item that has been checked in (added) to the library system.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|---|
| V_NAME (Version ID) | varchar(7) | char(7) | Contains the unique identifier of the version. |
| V_E_NAME (Item ID) | varchar(9) null | char(9) | Contains the unique numeric identifier for the version of the parent item. |
| V_CHECKIN_USER (Checked In By User) | varchar(32) null | varchar2(32) null | Contains the user name of the person who checked in the version. |
| V_CHECKIN_GRP (Checked In By Group) | varchar(32) null | varchar2(32) null | Contains the active group of the user who checked in the version (at the time the version was checked in). |
| V_CREATE_DATE (Checkin Date) | datetime null | date null | Contains the time and date when the version was checked in. |
| V_COMMENT (Comment) | varchar(224) null | varchar2(224) null | Contains a comment about the version. |
| V_APPLICATION (Application) | varchar(32) null | varchar2(32) null | Identifies the software application used to create the version. |
| V_FILE_DATE (Modification Date) | datetime null | date null | (Protected items only) Contains the time and date the version was last modified (with another application). |
| V_FILE_SIZE (File Size) | int null | numeric(10) null | (Protected items only) Indicates the size of the version (in bytes). |
| V_STACK_DATE (Storage Date) | datetime null | date null | (Internal use only) Contains the time and date the version was placed in the storage repository. |
| V_STACK_ID (Storage Repository) | int null | numeric(10) null | Contains the unique numeric identifier of the storage repository where the version is stored. |
| V_STACK_FILE_ID (Encrypted File Name) | varchar(12) null | varchar2(12) null | Contains the encrypted file name of the version as it is identified in the storage repository. |
| V_LAST_DATE (Last Access Date) | datetime null | date null | Contains the date and time that the version was last checked out, checked in, or copied. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|-----------------------|--|
| V_LAST_USER (<i>Last Access User</i>) | varchar(32) null | varchar2(32) null | Contains the user name of the person who last checked out, checked in, or copied the version. |
| V_LAST_GROUP (<i>Last Access Group</i>) | varchar(32) null | varchar2(32) null | Contains the name of the active group of the user identified in the V_LAST_USER property at the time the version was last accessed by that user |
| V_AVAIL_STAT (<i>Availability Status</i>) | int null | numeric(10) null | (Protected items only) Indicates the current storage status of the version. |
| V_DISP_ID (<i>Archive Repository</i>) | int null | numeric(10) null | (Internal use only) Contains the unique numeric identifier of the archive repository through which the version was moved offline. |
| V_DISP_DATE (<i>Archive Date</i>) | datetime null | date null | Contains the date and time the version was moved to the archive repository identified in the V_DISP_ID property. |
| V_RECLAIM (<i>Reclaim Pending?</i>) | int null | numeric(10) null | (Protected items only) Indicates whether the version has been flagged to be moved back into Content Services from offline storage. |
| V_PERMANENT (<i>Permanent Version?</i>) | int null | numeric(10) null | (Protected items only) Indicates whether the version should be exempt from automatic archival. |
| V_CHECKED_OUT (<i>Checked Out?</i>) | int null | numeric(10) null | Indicates the number of checkouts for this version. |
| V_MERGE_SRC (<i>Merge-Source Version ID</i>) | varchar(7) null | char2(7) null | (Versions on the primary version line only) Contains the unique identifier of a Version on the secondary version line if that version was merged into the primary version line to create this version. |
| V_MERGE_DST (<i>Merge-Destination Version ID</i>) | varchar(7) null | char2(7) null | (Versions on a secondary version line only) Contains the unique identifier of the version that was created if and when the secondary version line was merged into the primary version line. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|---------------------|----------------------|--|
| V_BRANCH_COUNT (<i>Secondary Version Line Count</i>) | int null | numeric(10) null | Contains the number of secondary version lines currently in the library system for this version. |
| V_BRANCH_CURVER (<i>Latest Version?</i>) | int null | numeric(10) null | (Internal use only) Indicates whether this is the most recent version. |
| V_REC_DATE (<i>Reclaim Submitted Date</i>) | datetime null | date null | (Protected items only) Contains the date and time when the Reclaim command was last used to submit a request to bring this version back online. |
| V_B_I_ID_MAJOR (<i>Full Backup ID</i>) | int null | numeric(10) null | (Internal use only) Contains the full backup portion of the backup ID for the most recent system backup in which the version was included. |
| V_B_I_ID_MINOR (<i>Incremental Backup ID</i>) | int null | numeric(10) null | (Internal use only) Contains the incremental backup portion of the backup ID for the most recent system backup in which the version was included. |
| V_CHECKOUT_DATE (<i>Source Version Checkout Date</i>) | datetime null | date null | Contains the date and time when the previous version was checked out to create this version. |
| V_OFFLINE_LOCATION (<i>Offline Location</i>) | varchar(32) null | varchar2(32) null | The offline location of the version (for example, archive tape number). |
| V_CSI_ID (<i>Content Search Repository</i>) | int null | numeric(10) null | Contains the unique ID assigned to the content search repository in which this version indexed. |
| V_CSI_CID (<i>Index ID</i>) | int null | numeric(10) null | Contains the unique ID that identifies this version in the content search repository. |
| V_CSI_STATUS (<i>Indexing Status</i>) | tinyint null | numeric(3) null | Indicates the current index status of the version in the content search repository. |
| V_CSI_MAJOR (<i>Index Full Backup ID</i>) | smallint null | numeric(5) null | (Internal use only) Contains the full-backup portion of the backup ID for the most recent system backup in which this version was included as an indexed version in the content search repository. |

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|--|-------------------------|-------------------------|---|
| V_CSI_MINOR (<i>Index Incremental Backup ID</i>) | smallint null | numeric(5) null | (Internal use only) Contains the incremental-backup portion of the backup ID for the most recent system backup in which this version was included as an indexed version in the content search repository. |
| V_CSI_PERMANENT (<i>Permanent Index?</i>) | tinyint null | numeric(3) null | Indicates whether the version should be exempt from automatic de-indexing. |
| V_FILE_NAME (<i>File Name</i>) | varchar(256) null | varchar2(256) null | Contains the file name for the version. The Oracle value also contains the file name extension. |
| V_SUBCLASS (<i>Reserved for future use</i>) | Reserved for future use | Reserved for future use | Reserved for future use |
| V_CD (<i>Child Count</i>) | int null | numeric(10) null | Specifies the number of child documents for this version if it is part of a compound document. |

VERSION_CHECKOUT Table

VERSION_CHECKOUT Table (SPI_OBJ_VERSION_CHECKOUT, 768)

Describes a particular checkout that was created with the spiVersionCheckout command.

| Property | MSSQL Type/Length | ORACLE Type/Length | Description |
|---|----------------------|------------------------|---|
| VC_E_NAME (<i>Item ID</i>) | varchar(9) | char(9) | Contains the unique numeric identifier of the version of the parent item. |
| VC_NAME (<i>Version ID</i>) | varchar(7) | char(7) | Contains the unique identifier of the version that was checked out. |
| VC_USER (<i>Checked Out By User</i>) | varchar(32) null | varchar2(32) null | Contains the name of the user who checked out the version. |
| VC_GROUP (<i>Checked Out By Group</i>) | varchar(32) null | varchar2(32) null | Contains the name of the active group when the version was checked out. |
| VC_PROTECTED (<i>Protected?</i>) | int null | numeric(10) null | Indicates whether the version is for a protected item. |
| VC_TYPE (<i>Same Version Line?</i>) | varchar(8) null | char(8) null | Indicates whether a check in will add a new version to the same version line from which the version is checked out. |
| VC_CHKOUT_DATE (<i>Checkout Date</i>) | datetime null | date null | Contains the date and time when the version was checked out. |
| VC_COMMENT (<i>Checkout Comment</i>) | text null | varchar2(2000) null | Contains a comment about the version. |
| VC_PATH (<i>Checkout Path</i>) | text null | varchar2(1024) null | Protected items only) Contains the path name (disk drive, directories, and file name) that identifies where the version was sent when it was checked out. |
| VC_WS_ADR (<i>Client Workstation Address</i>) | varchar(32) null | varchar2(32) null | (Protected items only) Contains the name that identifies the client machine where the version was sent when it was checked out. |

Sample SQL Queries

This section contains sample Structured Query Language (SQL) queries that show how to build an SQL query against a library system to obtain a data set. Using SQL to write reports against CS library system data is not difficult, provided you have a working knowledge of SQL and can connect to a library system database. Remember that while making that connection, the CS library system name corresponds to a database instance or tablespace. For example, the following samples were created using an CS library system named *undertow*; the SQL reporting utility (iSQL/w) had to connect to the associated database instance (in this case in Microsoft SQL Server) also named *undertow*.

“Sample Query 1” on page 570

“Sample Query 2” on page 571

“Sample Query 3” on page 572

Sample Query 1

Query:

```
select sa_alloc_name  
from storage_alloc
```

Result:

```
sa_alloc_name  
-----
```

Accounting

Documents

(2 rows affected)

Comment:

This query shows what storage categories have been created. The query uses the SELECT command to get all values of sa_alloc_name FROM the storage_alloc table.

[“Sample Query 2” on page 571](#)

[“Sample Query 3” on page 572](#)

Sample Query 2

Query:

```
select se_id, se_owner, se_ws_adr
from sessiontab
order by se_date
```

Result:

| se_id | se_owner | se_ws_adr |
|-------|------------|------------|
| ----- | ----- | ----- |
| 3608 | ReplDaemon | undertow |
| 3694 | Doug2 | ru4real |
| 3735 | Admin | rimson |
| 3751 | Doug | ru4real |
| 3754 | Admin | slipstream |
| 3761 | Admin | slipstream |

(6 rows affected)

Comment:

This query shows fields identifying who is logged on and includes their network address, ordered by date and time.

[“Sample Query 1” on page 570](#)

[“Sample Query 3” on page 572](#)

Sample Query 3

Query:

```
select e_org_filename  
from element  
where e_curver_cko=1
```

Result:

e_org_filename

0.MSG

0.WPD

1.WPD

2.WPD

boot.ini

7016.txt

july98.xls

(7 rows affected)

Comment:

This query shows the file names of those documents that are currently checked out.

[“Sample Query 1” on page 570](#)

[“Sample Query 2” on page 571](#)

Status Codes

All codes are unique irrespective of their severity levels. The severity level may be any one of the following:

| | |
|-----------------|----------------------------|
| SPI_SEV_SUCCESS | Successful completion |
| SPI_SEV_INFO | Informational error |
| SPI_SEV_WARNING | Failure, operational error |
| SPI_SEV_ERROR | Failure, integrity error |
| SPI_SEV_FATAL | System corrupt error |

Status Codes, Return Values, and Descriptions

The following table lists the Content Services API status codes. For each code, the hexadecimal, decimal, and return value is included.

Note Status codes 110 through 117 are not included in this table because they are not used by the Content Services API for the general library system. Likewise, error codes returned by the Client Library Files Setup program are not documented here. For those error codes, refer to “[Client Library Files Setup Error Messages](#)” on page 212.

In addition, a short description is provided for each code, which gives you information on what caused the error and, where appropriate, the possible response.

| Status Code | Hex | Decimal | Return Value | Description |
|--------------------|------|---------|----------------------------|--|
| SPI_SEV_SUCCESS, 0 | 0 | 0 | SPI_STS_SUCCESS | The operation is successful. |
| SPI_SEV_WARNING, 1 | 4001 | 16385 | SPI_STS_SYSTEM_UNAVAILABLE | The Property Manager is unavailable. |
| SPI_SEV_WARNING, 2 | 4002 | 16386 | SPI_STS_INVALID_LOGIN | The user name or password is invalid. |
| SPI_SEV_WARNING, 3 | 4003 | 16387 | SPI_STS_VERSION_MISMATCH | The configuration contains versions of the software that are incompatible. |
| SPI_SEV_WARNING, 4 | 4004 | 16388 | SPI_STS_LOGIN_LIMIT | The licensed number of concurrent open sessions has been reached. |
| SPI_SEV_WARNING, 5 | 4005 | 16389 | SPI_STS_NOT_MEMBER | The user is not a member of the specified group. |
| SPI_SEV_WARNING, 6 | 4006 | 16390 | SPI_STS_NOT_ADMIN | The user does not have Admin access rights to the object. |

| Status Code | Hex | Decimal | Return Value | Description |
|---------------------|------|---------|-----------------------------|--|
| SPI_SEV_WARNING, 7 | 4007 | 16391 | SPI_STS_MODIFY_FAILED | Changes have not been accepted. |
| SPI_SEV_WARNING, 8 | 4008 | 16392 | SPI_STS_NO_MEMORY | No free memory is available. |
| SPI_SEV_WARNING, 9 | 4009 | 16393 | SPI_STS_FAILURE | The operation has not succeeded. |
| SPI_SEV_WARNING, 10 | 400A | 16394 | SPI_STS_BAD_SCTX | The session context is invalid. |
| SPI_SEV_WARNING, 11 | 400B | 16395 | SPI_STS_NO_SUCH_ITEM | No such item has been found. |
| SPI_SEV_WARNING, 12 | 400C | 16396 | SPI_STS_NO_SUCH_VERSION | No such version has been found. |
| SPI_SEV_WARNING, 13 | 400D | 16397 | SPI_STS_STACK_ERROR | An error has been detected in the file-storage area. |
| SPI_SEV_WARNING, 14 | 400E | 16398 | SPI_STS_VARIANTS_DISABLED | Secondary versions cannot be created for any versions. |
| SPI_SEV_WARNING, 15 | 400F | 16399 | SPI_STS_CHECKED_OUT | The version is checked out. |
| SPI_SEV_WARNING, 16 | 4010 | 16400 | SPI_STS_BAD_ITEM | The item is corrupt. |
| SPI_SEV_WARNING, 17 | 4011 | 16401 | SPI_STS_BAD_VERSION | The version is corrupt. |
| SPI_SEV_WARNING, 18 | 4012 | 16402 | SPI_STS_NO_FILE_ACCESS | Access to the version file is not allowed. |
| SPI_SEV_WARNING, 19 | 4013 | 16403 | SPI_STS_INSUFFICIENT_ACCESS | The user has insufficient access rights. |
| SPI_SEV_WARNING, 20 | 4014 | 16404 | SPI_STS_NO_VARIANT_CKOUTS | Secondary versions cannot be created for the item. |
| SPI_SEV_WARNING, 21 | 4015 | 16405 | SPI_STS_FILE_EXISTS | The version file does exist. |
| SPI_SEV_WARNING, 22 | 4016 | 16406 | SPI_STS_NOT_CHECKED_OUT | The version is not checked out. |
| SPI_SEV_WARNING, 23 | 4017 | 16407 | SPI_STS_FILE_NOT_FOUND | The file was not found in the designated path. |

| Status Code | Hex | Decimal | Return Value | Description |
|------------------------|------|---------|----------------------------------|--|
| SPI_SEV_WARNING, 24 | 4018 | 16408 | SPI_STS_INVALID_ PASSWORD | The password is invalid. |
| SPI_SEV_WARNING, 25 | 4019 | 16409 | SPI_STS_INSUFFICIENT_ PRIV | The user has insufficient rights to perform the operation. |
| SPI_SEV_WARNING, 26 | 401A | 16410 | SPI_STS_UNPROTECTED | The item is unprotected. |
| SPI_SEV_WARNING, 27 | 401B | 16411 | SPI_STS_VERSION_LIMIT | The online version limit has been reached for the specified Item. |
| SPI_SEV_WARNING, 28 | 401C | 16412 | SPI_STS_NOT_BRANCH | The server is not a property server for this system. |
| SPI_SEV_WARNING, 29 | 401D | 16413 | SPI_STS_NOT_CURRENT | The version is not the latest version on the primary version line. |
| SPI_SEV_WARNING, 30 | 401E | 16414 | SPI_STS_NOT_ IMPLEMENTED | The specified option has not been implemented. |
| SPI_SEV_WARNING, 31 | 401F | 16415 | SPI_STS_UPDATE_ CONFLICT | A conflict has been found in the updating operation. |
| SPI_SEV_WARNING, 32 | 4020 | 16416 | SPI_STS_NOT_ PROTECTED | The item is not a protected item. |
| SPI_SEV_WARNING, 33 | 4021 | 16417 | SPI_STS_NOT_ UNPROTECTED | The item is not an unprotected item. |
| SPI_SEV_WARNING, 34 | 4022 | 16418 | SPI_STS_NOT_ONLINE | The version is not online. |
| SPI_SEV_WARNING, 35 | 4023 | 16419 | SPI_STS_NOT_OFFLINE | The version is not offline. |
| SPI_SEV_WARNING, 36 | 4024 | 16420 | SPI_STS_BAD_SHELF | The storage repository is corrupt. |
| SPI_SEV_WARNING, 37 | 4025 | 16421 | SPI_STS_BAD_ DISPOSITION | The archive repository is corrupt. |
| SPI_SEV_WARNING, 38 | 4026 | 16422 | SPI_STS_NOT_ RECLAIMABLE | The version is not reclaimable to online status. |
| SPI_SEV_WARNING, 39 | 4027 | 16423 | SPI_STS_MAINLINE_ CHECKED_OUT | The current version on the primary version line is checked out. |

| Status Code | Hex | Decimal | Return Value | Description |
|------------------------|------|---------|--------------------------|---|
| SPI_SEV_WARNING, 40 | 4028 | 16424 | SPI_STS_BRANCH_LIMIT | The limit to the number of secondary version lines has been reached. |
| SPI_SEV_WARNING, 41 | 4029 | 16425 | SPI_STS_END_OF_SET | The end of the value set has been reached. |
| SPI_SEV_WARNING, 42 | 402A | 16426 | SPI_STS_DUPLICATE_KEY | A duplicate key has been detected. |
| SPI_SEV_WARNING, 43 | 402B | 16427 | SPI_STS_PERMANENT | The version cannot be moved off line during automatic archival. |
| SPI_SEV_WARNING, 44 | 402C | 16428 | SPI_STS_VARIANT_REQUIRED | The version must be a secondary version. |
| SPI_SEV_WARNING, 45 | 402D | 16429 | SPI_STS_NO_SUCH_OBJECT | No such object can be found. |
| SPI_SEV_SUCCESS, 46 | 2E | 46 | SPI_STS_RECLAIM_PENDING | A reclaim request is pending on the version. |
| SPI_SEV_WARNING, 47 | 402F | 16431 | SPI_STS_INSTANCE_LIMIT | The instance limit has been reached. |
| SPI_SEV_WARNING, 48 | 4030 | 16432 | SPI_STS_UNSPECIFIED_REQ | A required entry has not been specified. |
| SPI_SEV_WARNING, 49 | 4031 | 16433 | SPI_STS_NULL_KEY | A NULL key has been detected. |
| SPI_SEV_WARNING, 50 | 4032 | 16434 | SPI_STS_BAD_CATEGORY | The storage category is not available. |
| SPI_SEV_WARNING, 51 | 4033 | 16435 | SPI_STS_ONLINE_ZERO | The online version limit is set to 0. |
| SPI_SEV_SUCCESS, 52 | 34 | 52 | SPI_STS_MORE_ROWS | More rows have been detected for the query. |
| SPI_SEV_SUCCESS, 53 | 35 | 53 | SPI_STS_AUTO_DISPOSE | Automatic archival is successful. You will never receive this message. It is for internal FileNet use only. |
| SPI_SEV_WARNING, 54 | 4036 | 16438 | SPI_STS_DIR_NOT_ALLOWED | The directory cannot be used. |
| SPI_SEV_WARNING, 55 | 4037 | 16439 | SPI_STS_PATH_NOT_FOUND | The path cannot be found. |

| Status Code | Hex | Decimal | Return Value | Description |
|------------------------|------|---------|-----------------------------|--|
| SPI_SEV_SUCCESS, 56 | 38 | 56 | SPI_STS_NOT_DELETED | The object has not been deleted. When checking in a version with the SPI_CHECKIN_KEEP flag not specified, the checkin succeeded but the file delete failed. |
| SPI_SEV_WARNING, 57 | 4039 | 16441 | SPI_STS_NETWORK_UNAVAILABLE | The network is unavailable or not operating. |
| SPI_SEV_WARNING, 58 | 403A | 16442 | SPI_STS_VARIANT_CHECKED_OUT | The secondary version is already checked out. |
| SPI_SEV_WARNING, 59 | 403B | 16443 | SPI_STS_NO_RESOURCE | The Storage Manager is out of the system resource. |
| SPI_SEV_WARNING, 60 | 403C | 16444 | SPI_STS_NO_SUCH_SERVER | The server is undefined. |
| SPI_SEV_WARNING, 61 | 403D | 16445 | SPI_STS_LOST_COMM | Communication with the Storage Manager has been lost. |
| SPI_SEV_WARNING, 62 | 403E | 16446 | SPI_STS_COMM_ERROR | A communication error with the Storage or Property Manager has been detected. This error may be displayed during login because: (1) a property server is not logged in to the LAN, (2) all Library System Services dll's are not in one of the LIBPATH directories specified in config.sys (e.g., in \os2\dll), (3) a library system database session cannot be established, and (4) under OS/2 2.x, there is no entry in the sql.ini file for the library system property server. |
| SPI_SEV_WARNING, 63 | 403F | 16447 | SPI_STS_USER_IO_ERROR | A read/write error has been detected on the user's file. |
| SPI_SEV_WARNING, 64 | 4040 | 16448 | SPI_STS_USER_DISK_FULL | No space is available on the user's disk. |

| Status Code | Hex | Decimal | Return Value | Description |
|------------------------|------|---------|-----------------------------|--|
| SPI_SEV_WARNING, 65 | 4041 | 16449 | SPI_STS_SHELF_FULL | No space is available in the storage repository. |
| SPI_SEV_WARNING, 66 | 4042 | 16450 | SPI_STS_SHELF_NOTFOUND | The storage repository could not be found on the server. |
| SPI_SEV_WARNING, 67 | 4043 | 16451 | SPI_STS_SHELF_FILE_NOTFOUND | The file could not be found in the storage repository. |
| SPI_SEV_WARNING, 68 | 4044 | 16452 | SPI_STS_SHELF_FILE_EXISTS | The file already exists in the storage repository. |
| SPI_SEV_WARNING, 69 | 4045 | 16453 | SPI_STS_SHELF_IO_ERROR | A read/write error has been detected on the stored file. |
| SPI_SEV_WARNING, 70 | 3046 | 16454 | SPI_STS_DEADLOCK_VICTIM | A required resource is unavailable due to deadlock. |
| SPI_SEV_WARNING, 71 | 4047 | 16455 | SPI_STS_MISSING_CONFIG | The configuration for this server is missing. |
| SPI_SEV_WARNING, 72 | 4048 | 16456 | SPI_STS_SHELF_SUSPENDED | The Storage Manager is suspended. |
| SPI_SEV_SUCCESS, 73 | 49 | 73 | SPI_STS_RECLAIM_WRONG_FILE | The identity of the reclaimed version file is questionable. During a reclaim request, the reclaim succeeds but the size of the reclaim file is different than the size of the archived file. |
| SPI_SEV_WARNING, 74 | 404A | 16458 | SPI_STS_VARIANT_CLOSED | The secondary version line has been merged into the primary version line. |
| SPI_SEV_WARNING, 75 | 404B | 16459 | SPI_STS_NO_SERVER_AVAILABLE | No Storage Managers are available. |
| SPI_SEV_WARNING, 76 | 404C | 16460 | SPI_STS_VARIANT_NOT_CURRENT | The secondary version required is not the latest version. |
| SPI_SEV_WARNING, 77 | 404D | 16461 | SPI_STS_CATALOG_FULL | The property repository or transaction log is full. |

| Status Code | Hex | Decimal | Return Value | Description |
|------------------------|------|---------|------------------------------------|--|
| SPI_SEV_WARNING, 78 | 404E | 16462 | SPI_STS_DB_ RECONNECT | The Property Manager has been successfully reconnected. |
| SPI_SEV_WARNING, 79 | 404F | 16463 | SPI_STS_ALREADY_ PENDING | A reclaim request is already pending on the version. |
| SPI_SEV_WARNING, 80 | 4050 | 16464 | SPI_STS_DEVICE_NOT_ FOUND | An unknown shelf-server device was specified. |
| SPI_SEV_SUCCESS, 81 | 51 | 81 | SPI_STS_BAD_DEFAULT_ GROUP | The user logging in is no longer a member of the default group specified in the User object. |
| SPI_SEV_WARNING, 82 | 4052 | 16466 | SPI_STS_FILE_ CREATION_ERROR | Creation of the new file could not be completed. |
| SPI_SEV_WARNING, 83 | 4053 | 16467 | SPI_STS_INVALID_ LICENSE | An invalid license number has been detected. |
| SPI_SEV_WARNING, 84 | 4054 | 16468 | SPI_STS_BAD_ PARAMETER | A bad parameter was passed to the function. |
| SPI_SEV_WARNING, 85 | 4055 | 16469 | SPI_STS_DBI_TEMP_ PATH | The temporary path could not be located. |
| SPI_SEV_WARNING, 86 | 4056 | 16470 | SPI_STS_MAX_VALUE | The maximum value has been reached. |
| SPI_SEV_WARNING, 87 | 4057 | 16471 | SPI_STS_BAD_QUERY_ LHND | The SPI_QUERY_LHND passed in was invalid. |
| SPI_SEV_WARNING, 88 | 4058 | 16472 | SPI_STS_TRUNCATED | The property value is too long for the destination and has been truncated. |
| SPI_SEV_WARNING, 89 | 4059 | 16473 | SPI_STS_USER_LIMIT_ EXCEEDED | The limit for registered users at this system was exceeded. |
| SPI_SEV_WARNING, 90 | 405A | 16474 | SPI_STS_DEVICE_NOT_ ALLOWED | An invalid device was specified on the server. |
| SPI_SEV_WARNING, 91 | 405B | 16475 | SPI_STS_NET_LOGIN_ NOT_ALLOWED | Logging in using the alternate user name is not allowed. |
| SPI_SEV_WARNING, 92 | 405C | 16476 | SPI_STS_CMDLN_PSWD_ NOT_ALLOWED | A user password cannot be entered during command-line login. |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-------------------------------|---|
| SPI_SEV_WARNING, 93 | 405D | 16477 | SPI_STS_DBI_ERROR | A general DBI error has been detected. |
| SPI_SEV_WARNING, 94 | 405E | 16478 | SPI_STS_OUTSTANDING_CHECKOUTS | The item could not be deleted due to outstanding version checkouts. |
| SPI_SEV_WARNING, 95 | 405F | 16479 | SPI_STS_NET_UNAME_NOT_FOUND | Network user name not found. |
| SPI_SEV_WARNING, 96 | 4060 | 16480 | SPI_STS_BUFFER_OVERFLOW | Internal buffer overflow has been detected. |
| SPI_SEV_WARNING, 97 | 4061 | 16481 | SPI_STS_LOGINS_DISABLED | Logons to the library system have been disabled. |
| SPI_SEV_WARNING, 98 | 4062 | 16482 | SPI_STS_REMOVAL_NOT_ALLOWED | A user cannot remove self from group. |
| SPI_SEV_WARNING, 99 | 4063 | 16483 | SPI_STS_DUPLICATE_NET_NAME | The user network name has already been used. |
| SPI_SEV_WARNING, 100 | 4064 | 16484 | SPI_STS_BAD_OAB_TYPE | Wrong OAB type for row function. |
| SPI_SEV_WARNING, 101 | 4065 | 16485 | SPI_STS_BAD_OBJ_TYPE | Inappropriate object type for function. |
| SPI_SEV_WARNING, 102 | 4066 | 16486 | SPI_STS_LOST_DB_CONNECTION | Connection to Property Manager has been lost. |
| SPI_SEV_WARNING, 103 | 4067 | 16487 | SPI_STS_FULL_BACKUP_PENDING | Full backup cannot be run because a previous full backup is pending. |
| SPI_SEV_WARNING, 104 | 4068 | 16488 | SPI_STS_KEYWORD_NOT_FOUND | Keyword for item could not be found. |
| SPI_SEV_WARNING, 105 | 4069 | 16489 | SPI_STS_MAXDBCLIENTS_EXCEEDED | Request would exceed database login limit. |
| SPI_SEV_WARNING, 106 | 406A | 16490 | SPI_STS_SHELF_EXISTS | The storage repository already exists. |
| SPI_SEV_WARNING, 107 | 406B | 16491 | SPI_STS_CATALOG_DISK_FULL | The backup failed because there is not enough disk space to create or process the dkupxxxx.yyy files (backup files of the property repository and transaction log). |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-----------------------------------|---|
| SPI_SEV_WARNING, 108 | 406C | 16492 | SPI_STS_SESSION_ ACTIVE | This session is active and, therefore, cannot be deleted. |
| SPI_SEV_WARNING, 109 | 406D | 16493 | SPI_STS_SESSION_ SUSPENDED | This session is already suspended. |
| SPI_SEV_WARNING, 118 | 4076 | 16502 | SPI_STS_SERVER_OUT_ OF_SERVICE | Operation of the Storage Manager is suspended. |
| SPI_SEV_WARNING, 119 | 4077 | 16503 | SPI_STS_DB_ INCONSISTENT | The Property Manager is in inconsistent state. |
| SPI_SEV_WARNING, 120 | 4078 | 16504 | SPI_STS_CONNECT_ REFUSED | Because of access problem, IPC could not be opened. |
| SPI_SEV_WARNING, 121 | 4079 | 16505 | SPI_STS_MAXQHNDSD_ EXCEEDED | Attempted allocation of MAXQHNDSD+1. |
| SPI_SEV_WARNING, 122 | 407A | 16506 | SPI_STS_STORED_PROC_ NOT_FOUND | Referenced stored proc not found. |
| SPI_SEV_WARNING, 123 | 407B | 16507 | SPI_STS_BACKUP_ DEVICE_FULL | (Currently not used.) |
| SPI_SEV_WARNING, 124 | 407C | 16508 | SPI_STS_USER_NOT_ ACTIVE | Since the user is not enabled, that user cannot log in to the library system. |
| SPI_SEV_WARNING, 125 | 407D | 16509 | SPI_STS_GROUP_NOT_ ACTIVE | The group is disabled (the user cannot change to this group). |
| SPI_SEV_WARNING, 126 | 407E | 16510 | SPI_STS_BACKUP_DB | (Currently not used.) |
| SPI_SEV_WARNING, 127 | 407F | 16511 | SPI_STS_BACKUP_RETRY | (Currently not used.) |
| SPI_SEV_WARNING, 128 | 4080 | 16512 | SPI_STS_BACKUP_ MUSTBEFULL | (Currently not used.) |
| SPI_SEV_WARNING, 129 | 4081 | 16513 | SPI_STS_BACKUP_ INTERRUPT | (Currently not used.) |
| SPI_SEV_WARNING, 130 | 4082 | 16514 | SPI_STS_BACKUP_ ASKSPACE | (Currently not used.). |
| SPI_SEV_WARNING, 131 | 4083 | 16515 | SPI_STS_BACKUP_ NOSPACE | (Currently not used.) |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-------------------------------------|--|
| SPI_SEV_SUCCESS, 132 | 84 | 132 | SPI_STS_DEFGROUP_ NOT_ACTIVE | Since the user's default group is disabled, the user was logged on as member of the General Users group. |
| SPI_SEV_WARNING, 133 | 4085 | 16517 | SPI_STS_NOT_IN_ SINGLE_USER_MODE | Since more than one user session is open, this delete/rename user/group operation cannot be completed. |
| SPI_SEV_WARNING, 134 | 4086 | 16518 | SPI_STS_BACKUP_ TRANFUL | (Currently not used.) |
| SPI_SEV_WARNING, 135 | 4087 | 16519 | SPI_STS_BACKUP_ TRUNCATE | (Currently not used.) |
| SPI_SEV_WARNING, 136 | 4088 | 16520 | SPI_STS_NO_SUCH_ USER | No such User object exists in the system. |
| SPI_SEV_WARNING, 137 | 4089 | 16521 | SPI_STS_NO_SUCH_ GROUP | No such Group object exists in the system. |
| SPI_SEV_WARNING, 138 | 408A | 16522 | SPI_STS_NO_SUCH_ SYSTEM | No such System object exists in the system. |
| SPI_SEV_WARNING, 139 | 408B | 16523 | SPI_STS_SERVER_ RESTRICTED | The Storage Manager is in a restricted state due to property repository or transaction log reaching full threshold. Clear log, increase property repository size, or change threshold. |
| SPI_SEV_WARNING, 140 | 408C | 16524 | SPI_STS_RESC_LIMIT_ EXCEEDED | The Property Manager database resource limit has been exceeded. |
| SPI_SEV_WARNING, 141 | 408D | 16525 | SPI_STS_CATALOG_ ERROR | An error has been detected in the ERRORLOG file (for example, meaning the SPI_SYS_LOGINS_ENABLED value cannot be set to Yes). |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-----------------------------|---|
| SPI_SEV_WARNING, 142 | 408E | 16526 | SPI_STS_NO_QRY | No properties have been specified for the search request in the SPI_SR_SHOW array. |
| SPI_SEV_WARNING, 143 | 408F | 16527 | SPI_STS_TEXT_SYNTAX_ERROR | A syntax error has been detected in the content search. |
| SPI_SEV_WARNING, 144 | 4090 | 16528 | SPI_STS_CSP_DB_ERROR | A database error has been detected in the content search engine. |
| SPI_SEV_WARNING, 145 | 4091 | 16529 | SPI_STS_MAX_SNIOPEN_RETRIES | The maximum number of retries has been reached while trying to establish network communications. The receiver was too busy to handle the request. |
| SPI_SEV_WARNING, 146 | 4092 | 16530 | SPI_STS_SNIOPEN_ERROR | An error has been detected while trying to establish the network communication using the sni_open call. |
| SPI_SEV_WARNING, 147 | 4093 | 16531 | SPI_STS_SNIMODE_ERROR | Issued for one of two reason: (a) after the initial sni_open call, an error is detected while attempting to set the state (byte mode) of the named pipe or socket or (b) after the state of the named pipe or socket is set, an error is detected during a check of that byte mode state. |
| SPI_SEV_WARNING, 148 | 4094 | 16532 | SPI_STS_SNIREAD_ERROR | An error has been detected while reading data during network communications. |
| SPI_SEV_WARNING, 149 | 4095 | 16533 | SPI_STS_SNIWRITE_ERROR | An error has been detected while writing data during network communications. |
| SPI_SEV_WARNING, 150 | 4096 | 16534 | SPI_STS_SNIXACT_ERROR | An error has been detected while performing a TRANSACT during network communications. |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|------------------------------------|--|
| SPI_SEV_WARNING, 151 | 4097 | 16535 | SPI_STS_SNICLOSE_ ERROR | An error has been detected while trying to close network communications. |
| SPI_SEV_WARNING, 152 | 4098 | 16536 | SPI_STS SOCK_HOST_ NOT_FOUND | An entry was not found in the hosts file. |
| SPI_SEV_WARNING, 153 | 4099 | 16537 | SPI_STS SOCK_SERVICE_ NOT_FOUND | An entry was not found in the services file. |
| SPI_SEV_WARNING, 154 | 409A | 16538 | SPI_STS SOCK_BIND_ ERROR | An error has been detected while trying to establish network communications. |
| SPI_SEV_WARNING, 155 | 409B | 16539 | SPI_STS_BACKUP_ TIMESYNC | (Currently not used.) |
| SPI_SEV_WARNING, 156 | 409C | 16540 | SPI_STS_BACKUP_ TIMEOUT | (Currently not used.) |
| SPI_SEV_WARNING, 157 | 409D | 16541 | SPI_STS_CSI_NO_FILE_ ACCESS | The user does not have access to the file for content search indexing. |
| SPI_SEV_WARNING, 158 | 409E | 16542 | SPI_STS_CSI_NO_INDEX | The content search repository cannot be accessed or does not exist. |
| SPI_SEV_WARNING, 159 | 409F | 16543 | SPI_STS_CSI_BAD_INDEX | The content search repository ID is invalid. |
| SPI_SEV_WARNING, 160 | 40A0 | 16544 | SPI_STS_CSI_LIMIT_ EXCEEDED | The item's indexed version limit has been exceeded. |
| SPI_SEV_WARNING, 161 | 40A1 | 16545 | SPI_STS_CSI_RESOURCE_ FAILURE | A content search repository resource failed. |
| SPI_SEV_WARNING, 162 | 40A2 | 16546 | SPI_STS_CSI_INDEX_ INCONSISTENT | An inconsistency has been detected in the content search repository. |
| SPI_SEV_WARNING, 163 | 40A3 | 16547 | SPI_STS_CSI_BAD_FILTER | The file filter is invalid. |
| SPI_SEV_WARNING, 164 | 40A4 | 16548 | SPI_STS_CSI_FAILURE | A general failure of the content search repository has been detected. |
| SPI_SEV_SUCCESS, 165 | A5 | 165 | SPI_STS_CSI_PARTIAL_ SEARCH | A partial content search succeeded. |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|--------------------------------|--|
| SPI_SEV_WARNING, 166 | 40A6 | 16550 | SPI_STS_CSI_DEADLK | A deadlock has occurred during the content search. |
| SPI_SEV_WARNING, 167 | 40A7 | 16551 | SPI_STS_DBI_INIT_FAILURE | The database library failed to initialize. For DOS, this failure typically indicates that the DBNMPPIPE, DBHPTCP, or DBWOLTCP TSR is not loaded. |
| SPI_SEV_WARNING, 168 | 40A8 | 16552 | SPI_STS_DSCONF_READ_FAILURE | A failure has occurred while reading the dsconf.ini file. |
| SPI_SEV_WARNING, 169 | 40A9 | 16553 | SPI_STS_NO_CSS_AVAILABLE | No Content Search Manager is available for the specified pipe. |
| SPI_SEV_WARNING, 170 | 40AA | 16554 | SPI_STS_BACKUP_VERIFY | (Currently not used.) |
| SPI_SEV_WARNING, 171 | 40AB | 16555 | SPI_STS_VERIFY_FAIL | (Currently not used.) |
| SPI_SEV_WARNING, 172 | 40AC | 16556 | SPI_STS_INDEXED | The version has already been indexed for content searches. |
| SPI_SEV_WARNING, 173 | 40AD | 16557 | SPI_STS_NOT_INDEXED | The version is not indexed for content searches; therefore, it cannot be deindexed or reindexed. |
| SPI_SEV_WARNING, 174 | 40AE | 16558 | SPI_STS_EMPTY_FILE_INDEX_ERROR | An empty version file cannot be indexed for content searches. |
| SPI_SEV_WARNING, 175 | 40AF | 16559 | SPI_STS_CSI_SRCH_TEMPERR | The search cannot be completed because of a conflict with the content search indexing process. Try the search again. |
| SPI_SEV_WARNING, 176 | 40B0 | 16560 | SPI_STS_BAD_DATETIME | An invalid datetime value has been detected. |
| SPI_SEV_WARNING, 177 | 40B1 | 16561 | SPI_STS_SHELF_NOT_RECLAIMABLE | The archive repository must be nonreclaimable to allow deletion of archived versions. |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|--------------------------------|---|
| SPI_SEV_WARNING, 178 | 40B2 | 16562 | SPI_STS_CSI_CREATE_INDEX | The attempt to add a content search repository failed. |
| SPI_SEV_WARNING, 179 | 40B3 | 16563 | SPI_STS_INVALID_DEVICE_TYPE | The device type is invalid for this operation. For example, you cannot assign shelves to floppy disk drives. |
| SPI_SEV_WARNING, 180 | 40B4 | 16564 | SPI_STS_SNI_SESSIONS_EXCEEDED | The NET BIOS session limit has been exceeded. |
| SPI_SEV_WARNING, 181 | 40B5 | 16565 | SPI_STS_CS_NOT_ALLOWED | The Count Only property in the Search object cannot be set to Yes if the query includes a content search structure. |
| SPI_SEV_WARNING, 182 | 40B6 | 16566 | SPI_STS_BACKUP_OUT_OF_SEQUENCE | (Currently not used.) |
| SPI_SEV_WARNING, 183 | 40B7 | 16567 | SPI_STS_INDEX_LOCKED | This operation cannot be completed because the content search repository is locked. |
| SPI_SEV_WARNING, 184 | 40B8 | 16568 | SPI_STS_INDEX_DISK_FULL | There is not enough free space on the specified disk to create/restore the content search repository or to index more files. Make more disk space available or add a new content search repository at a different location. |
| SPI_SEV_WARNIN, 185 | 40B9 | 16569 | SPI_STS_INDEX_FILE_FOUND | The content search repository cannot be created/restored because one or more program files for this index already exist on the specified disk. |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-----------------------------------|--|
| SPI_SEV_WARNING, 186 | 40BA | 16570 | SPI_STS_IDX_STATUS_ NOT_FOUND | The status of the content search repository could not be determined. This usually indicates the versions of the fultext.ftc file and the srs_ft library file are not compatible. |
| SPI_SEV_SUCCESS, 187 | BB | 187 | SPI_STS_NO_MORE_ ROWS | There are no more rows to return for this query. |
| SPI_SEV_WARNING, 188 | 40BC | 16572 | SPI_STS_SNI_BAD_ HOST_HANDLE | (Currently not used.) |
| SPI_SEV_WARNING, 189 | 40BD | 16573 | SPI_STS_SNI_BAD_ HOST_SPEC | (Currently not used.) |
| SPI_SEV_WARNING, 190 | 40BE | 16574 | SPI_STS_SNI_BAD_ FUNC_SEQU | (FileNet internal use only) |
| SPI_SEV_WARNING, 191 | 40BF | 16575 | SPI_STS_SNI_CON_ PROVDR_PROB | Named pipes communication could not be established. Typically, this means that named pipes is not functioning or permissions are set incorrectly. |
| SPI_SEV_WARNING, 192 | 40C0 | 16576 | SPI_STS_SNI_BUF_TOO_ SMALL | (FileNet internal use only) |
| SPI_SEV_WARNING, 193 | 40C1 | 16577 | SPI_STS_SNI_SYSTEM_ ERROR | (FileNet internal use only) |
| SPI_SEV_WARNING, 194 | 40C2 | 16578 | SPI_STS_SNI_NOT_ SUPPORTED | (FileNet internal use only) |
| SPI_SEV_WARNING, 195 | 40C3 | 16579 | SPI_STS_SNI_NO_ DATA_AVAILABLE | (FileNet internal use only) |
| SPI_SEV_WARNING, 196 | 40C4 | 16580 | SPI_STS_SNI_NO_ DATA_UNIT | (Currently not used.) |
| SPI_SEV_WARNING, 197 | 40C5 | 16581 | SPI_STS_SNI_FLOW_ CONTROL | (Currently not used.) |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|----------------------------------|---|
| SPI_SEV_WARNING, 198 | 40C6 | 16582 | SPI_STS_CAT_INDEX_ FULL | Because the maximum number of files have been indexed in the content search repository, pending indexing actions will be deleted; add new content search repository, reassign storage categories, or deindex some files. |
| SPI_SEV_WARNING, 199 | 40C7 | 16583 | SPI_STS_BACKUP_PATH | (Currently not used.) |
| SPI_SEV_WARNING, 200 | 40C8 | 16584 | SPI_STS_BK_NO_ CONNECT_EXCL_MODE | (Currently not used.) |
| SPI_SEV_WARNING, 201 | 40C9 | 16585 | SPI_STS_BK_IN_ EXCLUSIVE_MODE | (Currently not used.) |
| SPI_SEV_WARNING, 202 | 40CA | 16586 | SPI_STS_SCTX_IN_USE | Already using this SCTX in some thread. |
| SPI_SEV_WARNING, 203 | 40CB | 16587 | SPI_STS_ROLL_FORWARD | (DB2 only) The roll forward restore is pending. |
| SPI_SEV_SUCCESS, 204 | CC | 204 | SPI_STS_INDEX_NOT_ CREATED | The custom property was created successfully, however the index creation failed. |
| SPI_SEV_WARNING, 205 | 40CD | 16589 | SPI_STS_DB_NEEDS_ RECOVERY | This personal library system needs to use dsstrt to recover the Property Manager. The R:BASE database engine has had a transaction failure. Stop and restart the system (using dsstop and dsstrt) to allow users to once again connect to the system. |
| SPI_SEV_WARNING, 206 | 40CE | 16590 | SPI_STS_FULL_BACKUP_ FIRST | (Currently not used.) |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-----------------------------|---|
| SPI_SEV_SUCCESS, 207 | CF | 207 | SPI_STS_CSI_DISABLED | Content search indexing of files is currently disabled on this system (other portions of the operation, other than content search indexing/deindexing, completed successfully). |
| SPI_SEV_WARNING, 208 | 40D0 | 16592 | SPI_STS_BAD_ATTR_TYPE | You have specified an inappropriate property type for the function used. |
| SPI_SEV_WARNING, 209 | 40D1 | 16593 | SPI_STS_BAD_OPERATOR | The search operator is not valid for a property's data type. Check that a property with a Text data type is not using an operator other than LIKE or NOT LIKE (the only supported operators for Text data types). |
| SPI_SEV_WARNING, 210 | 40D2 | 16594 | SPI_STS_LOGICAL_PAGE_ERROR | (Currently not used.) |
| SPI_SEV_WARNING, 211 | 40D3 | 16595 | SPI_STS_RSTR_DB_TOO_SMALL | (Currently not used.) |
| SPI_SEV_WARNING, 212 | 40D4 | 16596 | SPI_STS_RSTR_DIFF_CHARSET | (Currently not used.) |
| SPI_SEV_WARNING, 213 | 40D5 | 16597 | SPI_STS_TEMPDB_FULL | (For Oracle databases) The temporary tablespace allocation is full. |
| SPI_SEV_WARNING, 214 | 40D6 | 16598 | SPI_STS_TAPE_DEVICE_ERROR | (Currently not used.) |
| SPI_SEV_WARNING, 215 | 40D7 | 16599 | SPI_STS_WRONG_TAPE_DEVICE | (Currently not used.) |
| SPI_SEV_WARNING, 216 | 40D8 | 16600 | SPI_STS_TAPE_TOO_SMALL | (Currently not used.) |
| SPI_SEV_WARNING, 217 | 40D9 | 16601 | SPI_STS_BACKUP_SERVER_ERROR | (Currently not used.) |
| SPI_SEV_WARNING, 218 | 40DA | 16602 | SPI_STS_BAD_PROPERTY_LEN | The custom property length is invalid. |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|-----------------------------------|---|
| SPI_SEV_WARNING, 219 | 40DB | 16603 | SPI_STS_SHELF_TIMEOUT | The Storage Manager was unable to process a request within its timeout period. |
| SPI_SEV_WARNING, 220 | 40DC | 16604 | SPI_STS_STRONG_REL | An object deletion failed because it was referenced by a strong relationship object. |
| SPI_SEV_WARNING, 221 | 40DD | 16605 | SPI_STS_DBI_TIMEOUT | (Currently not used.) |
| SPI_SEV_WARNING, 222 | 40DE | 16606 | SPI_STS_NESTED_TRANSACTION | A begin transaction was attempted within another transaction. |
| SPI_SEV_WARNING, 223 | 40DF | 16607 | SPI_STS_OBJECT_REF_ERROR | Operation cannot be completed because this object is being referenced by another object. |
| SPI_SEV_WARNING, 224 | 40E0 | 16608 | SPI_STS_USER_CONNECTIONS_EXCEEDED | There are insufficient database connections/processes allocated. |
| SPI_SEV_WARNING, 225 | 40E1 | 16609 | SPI_STS_MAX_EXTENTS_REACHED | <p>The Oracle extents maximum has been reached. This condition can be corrected by redefining the DEFAULT STORAGE CLAUSE for the library system tablespace to increase the MAX_EXTENTS setting (defaults to 210). The maximum value for MAX_EXTENTS per database object is 249.</p> <p>If the MAX_EXTENTS setting cannot be increased sufficiently, the SARO database objects for which the error was generated, must be recreated and restructured manually.</p> |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|------------------------------|---|
| SPI_SEV_WARNING, 226 | 40E2 | 16610 | SPI_STS_ACCESS_DENIED | Unable to access or create the specified directory. This is likely due to insufficient permissions to this directory. |
| SPI_SEV_WARNING, 227 | 40E3 | 16611 | SPI_STS_INCOMPATIBLE_DB | The wrong DLLs are installed for this database type. |
| SPI_SEV_WARNING, 228 | 40E4 | 16612 | SPI_STS_DBI_NO_RESULTS_ERROR | Indicates an internal database error where expected results were not returned. |
| SPI_SEV_WARNING, 229 | 40E5 | 16613 | SPI_STS_NO_SUCH_LANGUAGE | The specified language has not been supplied for this library system. |
| SPI_SEV_WARNING, 230 | 40E6 | 16614 | SPI_STS_WRONG_DLL | FileNet internal error code. |
| SPI_SEV_WARNING, 231 | 40E7 | 16615 | SPI_STS_INVALID_ORDER | Indicates an attempt to order by a list property (such as, a Keyword or Multivalue Custom Property). |
| SPI_SEV_WARNING, 232 | 40E8 | 16616 | SPI_STS_REPCHK_MAX_ERRORS | The maximum number of errors has been reached in a repository check. |
| SPI_SEV_WARNING, 233 | 40E9 | 16617 | SPI_STAT_LAST | You can define your own status codes by decrementing from MAXINT. Your defines should be compared against SPI_STAT_LAST to avoid collision. |
| SPI_SEV_WARNING, 232 | 40E8 | 16616 | SPI_STS_REPCHK_MAX_ERRORS | The maximum number of errors has been reached in a repository check. |
| SPI_SEV_WARNING, 233 | 40E9 | 16617 | SPI_STS_BROKEN_RELATIONSHIP | Indicates a weak relationship between objects was broken (e.g., between Version objects in a compound document). |

| Status Code | Hex | Decimal | Return Value | Description |
|-------------------------|------|---------|------------------------|--|
| SPI_SEV_WARNING, 234 | 40EA | 16618 | SPI_STS_NO_REPLICAS | Indicates a failed attempt to link to a replica (e.g., a replicated document). |
| SPI_SEV_SUCCESS, 235 | EB | 235 | SPI_STS_DYNAMIC_BROKEN | Indicates a dynamic link was successfully broken (i.e., that the modification of a Version object property has caused no version to satisfy a Dynamic property and Dynamic property value pair in a Dynamic relationship, which is a success condition). |
| SPI_SEV_WARNING, 236 | 40EC | 16620 | SPI_STS_PARENT_ONLINE | Indicates an archive action failed because the parent object is marked as being online rather than archived (e.g., for a compound document, a child document cannot be archived if its parent document is still online). |
| SPI_SEV_WARNING, 237 | 40ED | 16621 | SPI_STAT_LAST | You can define your own status codes by decrementing from MAXINT. Your defines should be compared against SPI_STAT_LAST to avoid collision. |

Replication Services



Replication Services supports many new features in this release, and this guide has been reorganized and rewritten extensively. Be sure to read this guide thoroughly before you implement Replication Services at your site. Replication Services can have profound impact on your library system configurations and the hardware and software resources required.

The Replication Services documentation is divided into these major topics:

- “Replication Services Overview” on page 594
- “Planning for Replication Services” on page 601
- “Replicating Objects” on page 609
- “Replication Home Server” on page 616
- “Setting Up Replication Services” on page 618
- “Enabling Replication Home Server” on page 625
- “Managing Replication Services” on page 629
- “Managing Replication Home Server” on page 644
- “Replication Recovery” on page 645
- “Replication Planning Worksheets” on page 652

Replication Services Overview


Replication Services gives your documents greater availability across your enterprise, automatically and without compromising document integrity. Instead of manually creating and maintaining the necessary user accounts and documents on separate libraries, replication automatically copies the documents and related security to another library system.

Replication Services provides a way to make documents and items in one library system available to other library systems. Replication Services automatically copies documents, items, and their security to additional specified library systems. As documents are checked out, updated, and checked in again, the changes are synchronized across the other libraries.

When a document is replicated, both the file content and the property values for standard properties and existing custom properties of the document are copied into the library systems participating in Replication Services. Users can quickly locate and access a replicated document anywhere in any replicated library. All instances of replicated documents or folders in any library system are periodically synchronized with any alterations to the original document or folder.

However, the synchronization is not immediate. That is, Replication is an asynchronous store-and-forward system. Changes are logged into a queue for processing, and then propagated to the replica libraries. The processing may take place almost immediately, or it may take place only during certain hours of the day, depending on your configuration parameters, the number and type of changes requested, the network load, and the availability of the replica libraries. For example, if a replica library is offline for maintenance (or a power failure), the update cannot be completed until the library is available again.

A replicated document's current status is reflected on all the participating library systems: if checked out, a document will be marked as checked out on all of these library systems. When a new version is checked in, that new version is replicated to all the participating library systems.



Replication Services Overview

Replication Services runs independently of the other library system services. Thus, Replication Services can continue running while other library system services are stopped and restarted and can automatically resynchronize with these library system services when they restart.

Using Replication Services can have a profound impact on your library system configurations and the hardware and software resources required. Take time to carefully plan your system before you implement replication in your network. Refer to [“Planning for Replication Services” on page 601](#) before you make any changes.

Notes

- For Replication to work correctly, each participating library system must have a unique name. Contact FileNet Professional Services if you have library systems with the same names and you need to implement replication.
- Replication Services is not a backup system. Use your database tools and file system backup tools to maintain proper backups of your libraries.

Replication Services Terminology

The following terms describe specific behaviors within Replication Services.

| Term | Meaning |
|-----------------------------------|--|
| Original Object | The original object that is replicated to other libraries. Modifications made to a original object are replicated to all replica objects automatically by the Replication Daemon. Original objects were previously called source objects. |
| Original Library | The library system where the original object is added. Original documents and original folders are created or added on the original library system. Original libraries were previously called source libraries. |
| Target Object Replica Object | The replicated (copied) object. Also referred to as a replica. For example, a folder replica or document replica. Modifications made to a target object are not always replicated back to the original object or any other replicated objects by the Replication Daemon. |
| Target Library Replica Library | The library system where replicas are located. The replica is created and owned by the Replication Daemon on the target library. |
| Replication Daemon | The replication software responsible for maintaining consistency between original objects and target objects. |
| ReplDaemon User | The Replication Daemon uses the ReplDaemon user to log in and out of the participating library systems to perform replication operations. The ReplDaemon user is assigned an encrypted password and has Administrator access rights to the specified library system. Do not modify this password. |
| Administrative Domain | <p>A collection of library systems that share security features. Within an administrative domain, any changes to User, Group, or Group Member objects are automatically replicated to the library systems. Stored searches created with the System Administration Tool, Document Classes, and CVLs are also replicated across the domain. The initial user and initial user group (Admin and Administrators group) and their associated stored searches are not replicated.</p> <p>Administrative domains are useful in environments where original documents may be located in a remote location and modifications are performed locally. Refer to “Administrative Domains” on page 606 and “Replicating Objects” on page 609 for more information.</p> |
| Domain Master | The library system where an administrative domain is created and maintained. For ease of maintenance, use the Domain Master to create and manage users, groups, Document Classes, and CVLs for the administrative domain. |
| Outside an Administrative Domain | All library systems that are not members of an administrative domain and that are participating in replication. Users can view, print, or copy replica documents. Users make all modifications on the original document, and all other libraries contain view-only document replicas. |

| Term | Meaning |
|---------------------------------|---|
| Inside an Administrative Domain | Replication features that apply only within an administrative domain: users can check out, check in, and modify any document replica (whose original document is inside the administrative domain) from any library system within the domain, provided that the user has the appropriate access rights. |
| Replication Folder | Refers to a root-level folder that the administrator uses to designate where documents and subfolders are replicated. Once a folder is replicated, documents and subfolders can be added from any participating library. Participating libraries are listed in the folder's Comment field. |
| Recovered Library | The damaged library system that is being synchronized. This library had some or all of its data destroyed due to a database or disk crash or failure. |
| Current Library | A library system participating in replication administrative domain that is used to recover the damaged library. Updated content information is retrieved from this library and replicated to the recovered library to bring it to the current state. |

Objects Supported for Replication

Replication Services always supports replicating the following objects between libraries:

- Folders
- Items
 - Documents
 - Compound Documents
 - Annotations for a Document
 - Published Documents
- Versions
- Custom property values when the same property exists on both libraries

Inside an administrative domain, these additional objects are replicated between the domain master and other libraries:

- Users
- Groups
- Group Memberships
- Searches stored at the user or group level creating using the System Administration Tools

- Default Item Access List
- Document Classes when those properties are identical on both libraries
- CVLs (Controlled Vocabulary Lists)

Administrative Domains

For purposes of replication, you can designate a collection of library systems as an *administrative domain*, which provides the following functionality:

- Users can check out and modify document replicas inside the administrative domain. Users can perform these operations from any library system within the domain (based on appropriate access rights).
- Users and groups are managed as a single security model. When you make changes to users or groups on their original library, the changes are replicated across all library systems within the domain. To simplify administration, migrate users and groups to the Domain Master using the USGRPMIG utility. The initial user and initial user group (Admin and Administrators group) and their associated stored searches are not replicated.
- Stored searches created with the System Administration Tools are replicated. Inside an administrative domain, Replication Services automatically replicates stored searches across all associated library systems for those properties that are identical on both libraries. If custom properties specified in a search do not exist on the target library, the search is not replicated.

Note Stored searches created with IDM Desktop are not replicated.

- Document Classes and CVLs are replicated within the domain provided that the custom properties exist on the target library system. Review the replication log for properties missing on the target library. It is the administrator's responsibility to ensure that the required custom properties are created on the target library.

Use the COPY_CP utility on the System Administration Tools CD to copy custom properties. Refer to [“COPY_CP” on page 382](#) for more information.

Administrative domains do not restrict replication of documents or folders to within a domain. You can replicate documents and folders to or from library systems in other administrative domains, or to library systems that are not members of an administrative domain. However, none of the domain-specific functionality applies outside of the domain.

For more information, refer to [“Administrative Domains” on page 606](#).

Administrative Domain Guidelines

- Users can modify replica documents inside an Administrative Domain. If your users only need access to view-only copies in the replica libraries, then you probably do not need to use an Administrative Domain.
- While users can checkout and modify documents from a replica library inside a domain, it is still best to add documents and folders only on the original library. See [Replicating Folder Structures](#) for more information.
- You can use an Administrative Domain to help synchronize users across library systems. If you have only a few users who need to participate in Replication Services, then you may have better results by not using an Administrative Domain. Instead, manually copy users from one library system to another.

Be very careful using tools like LDAP or DIRSYNCH within an Administrative Domain. Run these tools (or manually update users and groups) only on the Domain Master. Due to latency issues, making changes on more than one library will cause significant problems.

Outside an Administrative Domain

Documents replicated to library systems outside of an administrative domain are view-only documents. All modifications must be performed on the original document. The replicas can be copied, viewed, or printed.

The Replication Daemon

When you install Content Services, a ReplDaemon user is automatically created. The Replication Daemon uses the ReplDaemon user to log in and out of the participating library systems to perform replication operations. The ReplDaemon user is assigned an encrypted password and has Administrator access rights to the specified library system. Do **not** modify this password.

The Replication Daemon maintains consistency between the original library and the target libraries. It does this by creating work requests that contain add, delete, or update information for a specific object.

For example, when an object is selected for replication, a work request is created by the Replication Daemon on the original library system for the Replication Daemons on the target library systems. The target Replication Daemon then logs into the object's original library system (as the ReplDaemon user) and performs the work request that adds, deletes, or updates the replica object.

Note To synchronize Replication Services among the participating library systems, it is mandatory that all ReplDaemon users have Administrator privileges to the library systems and that their passwords are not changed. For this reason, the password is automatically assigned during the server installation and is algorithmically derived and encrypted.

The Replication Home Server feature

You can use the Replication Home Server feature to ensure centralized management of documents and folders from the Domain Master. You can implement this feature on one or more replication folders on the Domain Master. On implementing this feature, all the subfolders and documents within a replication folder are automatically sourced on the Domain Master. Thereby you can ensure that all management tasks related to these subfolders and documents are possible only from the Domain Master.

Once you enable Replication Home Server, you can perform the following operations only from the Domain Master and not from any member library:

- Delete documents or subfolders within the replication folder.
- Modify access rights for the documents or subfolders within the replication folder.

For more information, refer to [“Replication Home Server” on page 616](#).

Managing Document Classes and CVLs

Document Classes and Controlled Vocabulary Lists (CVLs) are only replicated inside an administrative domain. Outside an administrative domain, you must copy Document Classes and CVLs to each library using the import and export tools available in CS Explorer.

Keep the followings notes in mind when working with Document Classes and CVLs in a replication environment:

- Create and maintain the Document Classes and CVLs on one library system. For best results, use the Domain Master. When adding existing libraries into a replication domain, use the Document Class and CVL migration tool to migrate administrative functions to the Domain Master.
- Use the CS Document Class and CVL Admin to make changes.
- When you makes changes with the CS Document Class and CVL Admin, you need to close the tool and restart it to see the changes.
- After making changes, wait for a few replication cycles to complete for the changes to be replicated to the other library systems. Any users with IDM Desktop who were logged on while you made changes will need to log out and log back into the library to see your changes.
- CVL datatypes are validated during replication. If CVLs on library systems A and B have the same name but different datatypes, the replication is canceled. If the datatypes are the same, the corresponding CVLData objects will merge. When replicating a CVLDATA object from library system A to library system B, on which there is a CVL whose name matches the parent on A (there should be) and the datatypes differ, the replication from A to B is canceled.
- You cannot delete a Document Class that is currently used by existing documents.
- Do not modify the privileges for the Document Class custom property. Only the Owner Privileges should be set to Read/Write: no one else should have write privileges.
- When you add a new library to an existing administrative domain, Replication Services uses the following logic to resolve duplicate Document Class and CVL definitions:
 - a. Document Classes and CVLs on the Domain Master are copied to the new library system. This includes Document Classes and CVLs that originate on the Domain Master and on the existing domain member libraries. The Domain Master objects overwrite any duplicates on the new library.
 - b. Unique Document Classes and CVLs that originate on the new library will be replicated to the other library systems within the domain during the next replication cycle.

Planning for Replication Services

Before you set up Replication Services, you need to carefully define your business requirements and verify that Replication Services solves the business problem. If so, make sure your library system servers, the Content Services Admin Tools, and the IDM Desktop clients are set up appropriately. In general, you will have to determine which objects to replicate and how to administer user access.

You need to consider the following:

- The scope of Replication Services you need.
- The configuration of existing library systems.
- The configuration of existing Content Services Admin Tools and IDM Desktop clients.
- The intended distribution of replicated documents.
- The impact on your existing network infrastructure and hardware.
- The impact on your current backup and recovery scheme.
- The nature of the library systems you will establish.

Careful planning is essential before configuring and using Replication Services successfully. Some options, such as administrative domains, cannot be changed once you configure them. Before enabling and setting up Replication Services, you should carefully consider the issues raised in the following topics. Refer to the [“Configuration Worksheets” on page 43](#) that you used to design your library systems to help answer these questions. Additional worksheets are provided for planning Replication Services. Refer to [“Replication Planning Worksheets” on page 652](#) for more information.

Note Contact FileNet Technical Support if you need assistance in implementing Replication Services at your site.

The Scope of Replication Services Needed

To define the scope of Replication Services needed for your enterprise, you need to carefully consider the purpose of replicating documents at your site and how your users will interact with the replicated documents. Use the following questions and the [“Replication Planning Worksheets” on page 652](#) to determine your needs.

- What is the reason for making duplicate documents available across library systems?
- Which sites and library systems should participate in Replication Services?
- Which documents or kinds of documents should be replicated?
- Will users need to modify replicated documents?
- Will users only need to view, print, or copy documents?
- Which users and groups should be able to replicate documents?
- Do the users and groups exist on the libraries that will participate in replication?
- What replication Administrative Domains should be established?

- On which replication folders, do you need to implement the Replication Home Server feature?

Existing Library Systems

Replication Services will affect your existing library systems. You may need to upgrade your system software and configurations, and your end users may need to change how they work with documents within your site.

- The participating library systems must use the same library system version. If your site uses a software patch, the same patch must be installed on all participating libraries.
- The participating library systems must use the same database type and release version of MSSQL or Oracle. FileNet does not support replicating to a different database type.
- Replication Services requires that each participating library system has a unique name. This is necessary to differentiate between library systems for replicas. Contact FileNet Professional Services if your enterprise currently has multiple library systems with the same name and you plan to implement replication.
- Replication Services will change the way your users understand the document management architecture at your site. The existing design of library and folder trees, for example, will probably be altered as authorized users and groups can add documents and subfolders to existing replicated folders.

When a folder is replicated, all its subfolders and documents are replicated. A user with access rights to a top level folder may not have access to a subfolder, but the entire structure is replicated and all subfolders and documents may be viewed based on their security.

- You may need to reconfigure existing users and groups across library systems. For example, you may have duplicate group names with different users in different library systems. Administrative domains provide a domain-wide method for administering groups and users across library systems. If you do not use administrative domains, you will need to administer changes manually to duplicate groups and users on each library system. Refer to [“Administrative Domains” on page 606](#) for more information.
- The users or groups adding documents to replicated folders must understand that the replicated document may be accessed by different users on the target library systems.

The Intended Distribution of Replicated Documents

Once you determine the items to be replicated, you need to control their distribution. These factors affect distribution:

- The Replication folders you establish.

You can designate multiple Replication folders, each replicating documents to a subset of the participating library systems.

- The access rights you grant users and groups to each of these Replication folders.

For example, you may want one set of authorized users and groups to be able to replicate documents only to the library systems referenced by Replication folder A, and another set of authorized users and groups to be able to replicate documents only to the library systems referenced by Replication folder B.

- The replication administrative domains you set up.

You can establish an administrative domain to ensure that users can access and modify any document on a library within a particular domain. Or, you can restrict replica access to viewing and printing by replicating to

libraries outside an administrative domain or by not using administrative domains at all. Refer to [“Administrative Domains”](#) on page 606 for more information.

- The users and groups that exist inside an administrative domain and across the participating library systems.

Inside an administrative domain, Replication Services automatically replicates access list entries (Access Control object instances) between library systems, enabling you to manage security for all library systems within the domain as a single unit.

If you don't use administrative domains, or if documents are replicated outside of an administrative domain, Replication Services will not affect the existing security. You may need to create duplicate accounts on library systems outside the domain to provide view access to all intended users.

Impact on Network and Hardware Resources

Replication Services increases network traffic. Network outages and server down time affect synchronization across the participating library systems. For best results, you need a high-speed, reliable, and efficient network.

Replicas are copies of original documents and folders. Because users may access document replicas from different library systems and associated storage servers, you may need to change the following for existing library systems:

- Current load balancing required for your network.

Large numbers of documents being replicated during peak hours may place an additional load on your existing network resources. You can adjust the Replication Services configuration settings to control how frequently the Replication Daemon updates replicas in the system.

Documents that are expected to be modified frequently should probably not be replicated for purposes other than read-only: synchronization of the replicas across all of the library systems is not simultaneous. These operations require a minimum of one replication cycle to complete. For more information, refer to [“Replication Services Configuration Parameters”](#) on page 624.

- Storage space demands on library system servers participating in Replication Services.

Due to the increased number of transactions on your property servers, you may need to double the disk space allocated to transaction logs for MSSQL databases.

Backup and Restore Operations

Replicas can be backed up and restored in the same way as all other library system objects and document content.

However, to maintain correct synchronization of replicas and their original objects under all circumstances (such as failure of a library system's host server), make sure you schedule backup operations between library systems participating in replication at the same time. Doing so means you will be able to synchronize restore operations on all these library systems, and thereby maintain inter-library data integrity.

If you do not synchronize backups between library systems, the following can occur if it becomes necessary to restore an administrative domain from backups:

- Replication Services may not be able to correctly resynchronize the replicas with their original objects (at least until the original object is again modified).
- Objects replicated since the last backup may be lost (at least until the original object is again modified).

This is a by-product of the way the Replication Daemon marks and queues objects for replication. When a user modifies an original object, the program creates a work request for each target library. When Replication Services successfully replicates that object, it marks the work request as completed. If you subsequently restore the target library system to a state before the object was replicated, the original library system has no way of knowing which objects might be missing on the target library system. Only when/if a user modifies the object again, will Replication Services generate a new work request and cause the object to be replicated again on the target library systems, thereby eliminating the discrepancy.

Note Replication Services is *not* a backup system. Use your database tools and file system backup tools to maintain proper backups of your libraries.

The Original Library System

Replication Services is based on the original object: for any replica, properties and modifications originate in the original object. This has particular effect on those objects that may be different between library systems:

- Users and groups.

Outside an administrative domain, the same users and groups generally do not match between library systems. The Replication Services configuration options `ReplicateItemACLs` and `ReplicateFolderACLs` specify how to define a document or folder replica's user and group access list entries. For more information, refer to [“Replication Services Configuration Parameters” on page 624](#).

- Document or folder titles.

Outside a replication administrative domain, users can modify only original documents or folders. To help your end users identify the original library for a document or folder, you can include an indicator in the document or folder title. Refer to [“Replication Services Configuration Parameters” on page 624](#) for more information on the `ItemTitleFormat` and `FolderTitleFormat` settings.

- Custom properties.

Replication Services does not automatically replicate custom property definitions between library systems, either inside or outside an administrative domain. However, if a custom property is defined, Replication Services will replicate the value of the custom property whenever a user adds or checks in the document.

We suggest you maintain a consistent set of custom properties across the participating library systems. To do so, you must add any necessary custom properties to the library systems before enabling Replication Services.

Use the `COPY_CP` utility on the System Administration Tools CD to copy custom properties. Refer to [“COPY_CP” on page 382](#) for more information.

If needed, you can rename custom properties using the System Administration Tools. Remember to change the display name for the property also.

Replication of custom property values requires that all the following criteria be identical:

- Property name
- Data type
- Length
- Parent object

- Multivalue setting

Administrative Domains

You can designate a collection of library systems as an *administrative domain*, which provides the following functionality:


- Original and replicated document versions can be checked out and modified. Users can check out, cancel checkout of, and check in documents inside the administrative domain. During the check in process, changes to the replica version are copied to the original and to all replicas within the domain. Users can perform these operations from any library system within the domain (based on appropriate access rights).

Note You must perform the check in or cancel check out operation on the same library system where you performed the check out.

- Users and groups are managed as a single security model. Replication Services automatically replicates the following between library systems inside an administrative domain:
 - Users
 - Groups
 - Modifications to users (except the Admin user)
 - Modifications to groups (except the Admin group)
 - Modifications to group membership

Refer to [“Managing Users and Groups” on page 633](#) for more information.

- Document Classes and CVLs are replicated within the domain. Refer to [“Replicating Objects” on page 609](#) for more information.
- Stored searches created with the System Administration Tools are replicated within the domain. Refer to [“Replicating Objects” on page 609](#) for more information.
- The Domain Master is the library system where the administrative domain is created and maintained.



Administrative domains do not restrict replication of documents or folders to libraries within a domain. You can replicate documents and folders to or from library systems in other administrative domains, or to library systems that are not members of an administrative domain. However, none of the domain-specific functionality applies outside of the domain.

You can configure more than one administrative domain, but a library system can be a member of only one administrative domain.

Caution These restrictions apply to replication administrative domains:

- Once assigned to an administrative domain, a library system cannot be removed from that domain. Plan your domains carefully before you configure replication. Refer to “[Planning for Replication Services](#)” on page 601.
- A library system can be a member of only one administrative domain.
- If a duplicate user or group with the same name exists when you create an administrative domain or add libraries to an existing administrative domain, Replication Services designates the object with the earliest creation date as the original object and the remaining objects are overwritten. Review the replication log (**repl.log**) to identify how name conflicts were resolved.

For example, if a user named Danika exists on two library systems inside an administrative domain, the Danika user with the earliest creation date is designated the original User and is replicated to the other libraries. Check for any potential naming conflicts before configuring an administrative domain.

- If a duplicate DocClass or CVL with the same name exists when you create an administrative domain or add libraries to an existing administrative domain, the following logic is used to resolve the duplicate items:

- Document Classes and CVLs on the Domain Master are copied to the new library system. This includes Document Classes and CVLs that originate on the Domain Master and on the existing domain member libraries. The Domain Master objects overwrite any duplicates on the new library.
- If you added multiple libraries at the same time with duplicate Document Class or CVL names, then an error message is written to the replication log (**repl.log**) and the administrator must delete one of the objects from its library system. Review the replication log to identify name conflicts, and take appropriate action.

For ease of management, create all DocClasses and CVLs on the Domain Master before configuring the administrative domain. Use the DocClass/CVL Tool export feature to create export files for importing on the Domain Master. For more information, refer to the FileNet Content Services Admin Tools documentation.

Refer to “[Establishing a Replication Administrative Domain](#)” on page 621 for instructions on defining an administrative domain.

Replicating Objects

The objects that are replicated depend on whether you are replicating inside or outside of an administrative domain. Documents and subfolders within an administrative domain can be replicated to library systems outside a domain, but none of the domain features are supported.

| Outside an Administrative Domain | Inside an Administrative Domain |
|--|--|
| Folders | Same as Outside a Domain, plus: |
| Items | User |
| <ul style="list-style-type: none"> Documents Compound Documents Annotations for a Document Published Documents | Group |
| | Group Membership |
| | Default Item Access List |
| | Access Control List |
| Versions | Document Classes when properties are identical on both libraries |
| Custom Property values, if the same property exists on both libraries | CVLs (Controlled Vocabulary Lists) |
| | Search stored at the user or group profile created using the System Administration Tools |


Replicating Documents and Folders

Documents and subfolders added to replication folders are replicated between libraries. Documents and subfolders can be added from any library system.

You use CS Explorer to configure the replication folder, but other client interfaces (such as IDM Desktop) display the replication folders you create. Replication Services then replicates the designated replication folder to all of the library systems specified. Each folder or document that is replicated retains its original library identity on all of the participating library systems. IDM Desktop displays the Replication folders with this icon:

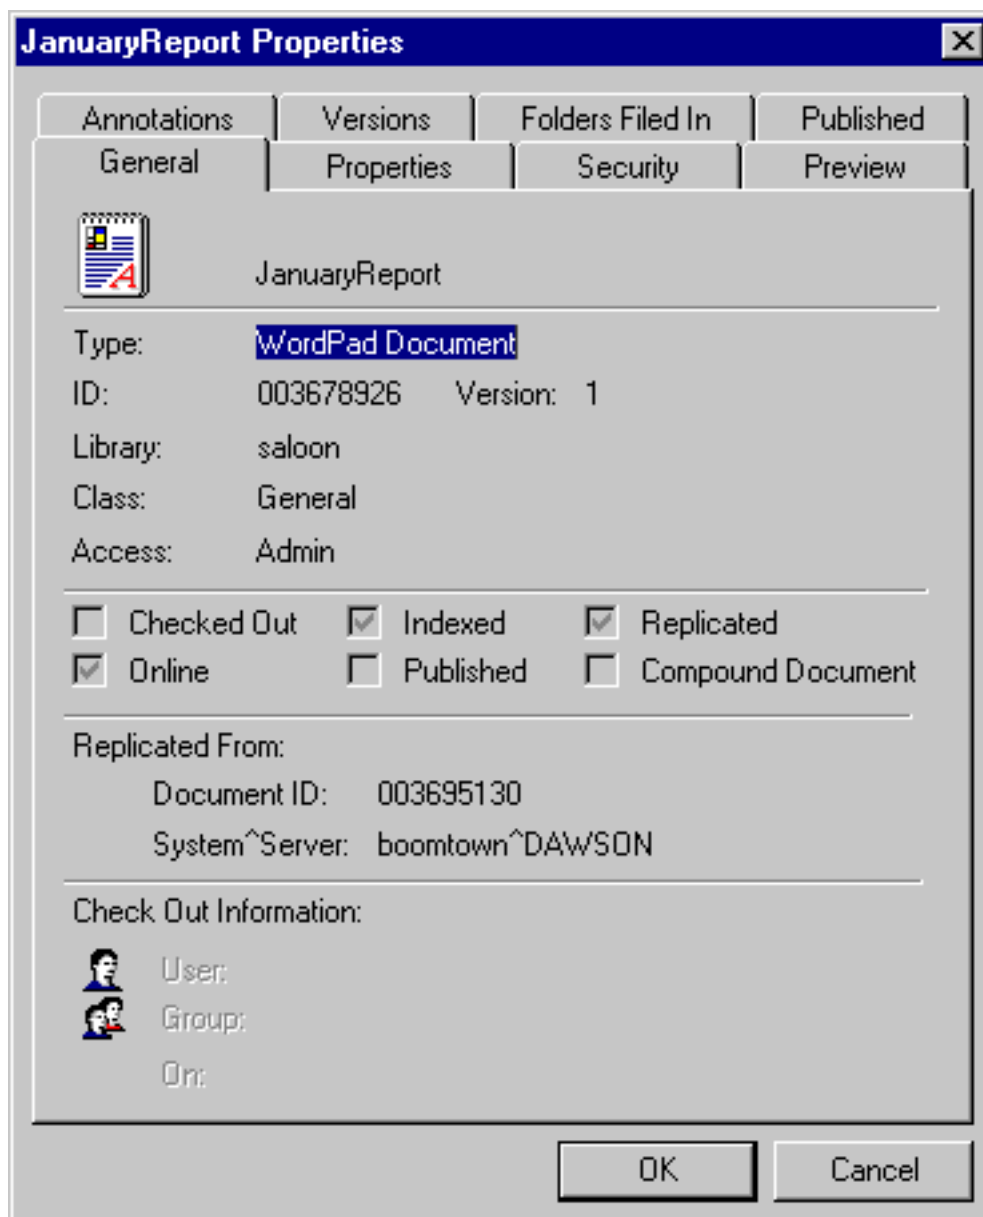


IDM Desktop displays replication information for replication folders and documents. In the following example, the replication folder named Proposals exists on the library named Saloon and replicates to the library named Zipper:



Note that the ID for the Proposals folder in the Saloon library is 94294465. You can determine if a folder is a replica or an original from its ID. An original replication folder or subfolder has a numeric ID. A replicated folder's ID begins with R, such as R38440d9a.

In the following example, the document named JanuaryReports exists on the library named Saloon as a replicated document. It is replicated from the library named Boomtown:



All of the objects (subfolders and documents) existing or subsequently placed in the replication folder are also replicated to the other library systems. For folders, this may be a single folder or a hierarchy (tree) of folders. A replication folder's tree structure is also retained across the participating library systems. Users must have Author access to the replication folder to add documents or folders for replication. It is possible to file a document in the target replication folder, and Replication Services will replicate it to the other libraries.

While users can check out and modify documents from a replica library inside a domain, it is still best to add documents and folders only to the original library.

There can be one or more replication folders for a library system, with each replication folder referencing different library systems.

Note If you attempt to copy or drag a replicated folder from one Replication folder tree to another, an error message will be returned.

Replicating Compound Documents

Compound documents are replicated both inside and outside of an administrative domain. When a user places a compound document into the replication folder, the compound document is replicated to other libraries. The links between the parent document and any child documents are maintained in the replica.

For best results, follow these guidelines:

- Keep the parent document and all of its child documents in the same folder in the same original library.

Note In case the child documents for a compound document are not kept in the replicated folder, the links between parent and child documents will not be maintained in the replica.

- Make all changes on the original library (even inside an Administrative Domain).

It is safe to check out and modify the content of a parent or child document from any library where it appears, but changes to the linkage between documents (including OLE embedding) should only be made on the original library. Since some types of links may not stand out upon casual inspection, it is safest to make all changes on the original library.

Note You cannot use a replicated document to create a compound document, and you cannot change the links within a replica of a compound document. You can only make these changes in the original document.

Replicating Published Documents

Published Documents are replicated both inside and outside of an administrative domain. When a user places a published document into the replication folder, the published document is replicated to other libraries.

The relationship between a published document and its original document is not replicated. That is, the published document and its original document are considered as separate unique documents. To replicate both documents, you must place both documents into the replication folder.

Replicating Document Annotations

Documents and their associated annotations are replicated both inside and outside of an administrative domain. When a user places an annotated document into the replication folder, the annotations are replicated with the document.

Note You cannot change a replica's annotation links. You can only make these changes in the original document.

Replicating Document Classes and CVLs

Document Classes and Controlled Vocabulary Lists (CVLs) are only replicated inside an administrative domain.

When you add or change a Document Class or CVL, the changes are replicated within the domain. If a specific custom property required by a Document Class does not exist on a target library system, that property is not replicated to that library system. Replication Services logs a warning that the replication of the Document Class is incomplete, and that it will attempt to resolve it at a later time. Review the replication log (**repl.log**) to determine which custom property is missing or not identical on the target library.

You must change a Document Class or CVL definition on the original library. For ease of management, create all Document Classes and CVLs on the Domain Master. Before creating an administrative domain with existing libraries, use the DocClass/CVL Tool export feature to create export files for importing on the Domain Master. For more information, refer to the Content Services Admin Tools documentation.

Use the Document Class and CVL migration tool to administer Document Classes and CVLs from the Domain Master.

When you add a new library to an existing administrative domain, Replication Services uses the following logic to resolve duplicate Document Class and CVL definitions:

1. Document Classes on the Domain Master are copied to the new library system. This includes Document Classes that originate on the Domain Master and on the existing domain member libraries. The Domain Master objects overwrite any duplicate Document Classes on the new library.
2. CVLs on the Domain Master are copied to the new library system. This includes CVLs that originate on the Domain Master and on the existing domain member libraries. The Domain Master objects merge any duplicate CVLs regardless of audiotape or values. As a result, you may have CVL data that originates on different libraries.
3. Unique Document Classes and CVLs that originate on the new library will be replicated to the other library systems within the domain during the next replication cycle.

Do not add multiple libraries at the same time. Error messages are written to the replication log (**repl.log**), and the administrator must delete one of the objects from its library system. You cannot delete Document Classes that are currently used by existing documents. Review the replication log to identify name conflicts.

Outside an administrative domain, you must copy Document Classes and CVLs to each library using the import and export tools available in CS Explorer. For more information, refer to the Content Services Admin Tools documentation.

A replicated document can reference a Document Class or CVL that does not exist on the target library system. IDM Desktop will display the default Document Class for that library system to the document.

For more information, see [“Managing Document Classes and CVLs” on page 600](#).

Replicating Custom Property Values

Custom property values are replicated both inside and outside of an administrative domain. Custom property values are replicated only if the same properties already exist on the target library systems.

The Replication Daemon attempts to locate the custom property on the library system. The property name, data type, length, parent object, and multivalue setting must be identical for the values to be replicated. If a custom property match is not found, the value of the custom property is not replicated with the document.

You need to set up and maintain a consistent set of custom property values across the participating library systems. To ensure consistency from the start, you must add any necessary custom property to the target library systems before enabling Replication Services.

Use the COPY_CP utility on the System Administration Tools CD to copy custom properties. Refer to [“COPY_CP” on page 382](#) for more information.

Note The custom properties defined on a target library affect how document properties, stored searches, Document Classes, and CVLs are replicated.

Replicating Users and Groups

Users and groups are only replicated inside an administrative domain.

When you add or modify a User, Group, Group Member, or Access Control object on any library system, Replication Services automatically replicates those objects across all participating library systems. The initial user and initial group (Admin and Administrators group) are not replicated. However, a user who is a member of the Administrators group will be replicated and group membership on the target will not include the initial group (Administrators group).

You should check for potential user and group naming conflicts before configuring an administrative domain. When more than one user or group object of the same name exists inside an administrative domain, Replication Services designates the object with the earliest creation date as the original object and the remaining objects as replicas. For example, if a user named Danika exists on two library systems inside an administrative domain, the Danika user with the earliest creation date will be designated the original User object instance, and the other Danika user will be designated a replica of this object instance.

Use the user/group migration utility USGRPMIG to migrate administration of users and groups to the Domain Master.

Caution Be sure to check that user and group names existing in different library systems actually belong to the same individuals. Otherwise, the Access List and the Default Item Access List may assign inappropriate access rights during replication to particular user or group replicas. For example, Replication Services could inadvertently assign user or group replicas viewer rights to documents and folders they should not be able to view.

For best results, you should use the Domain Master library system to add new users and groups to the domain. You can add new users and groups to the administrative domain from any library system within the administrative domain.

Inside an administrative domain, the Access Control object instance for a user or group is also replicated, with maximum access rights set to Author.

Outside of an administrative domain, Replication Services does not replicate users and groups. You must add such objects manually to all participating library systems, making sure to duplicate associated properties as appropriate for your site.

Replicating Stored Searches

Stored searches created with the System Administration Tools are replicated within an administrative domain.

All stored searches are automatically replicated to the other participating library systems if the custom properties used in the search already exist on the target library systems. The property name, data type, length, parent object, and multivalue setting must be identical for the values to be replicated. If a custom property match is not found, the search criteria is not replicated.

Users can change the Comments on a replica of a stored search, but the changes are not replicated to the original or other target libraries.

Note The custom properties defined on a target library affect how stored searches are replicated. Inside an administrative domain, Replication Services automatically replicates searches if all custom property values used in the search already exist on the target library.

Replicating Document Item and Version Properties

You can make changes to item and version properties for original documents. However, you must check out and check in the original document from the original library to replicate your changes to the respective target libraries. If you do not check in your changes, they are not replicated to the target libraries until the document is checked in again.

For document replicas, you make changes to the version properties during check in. Item properties cannot be changed on document replicas. Any changes you make to item properties on a replica are temporary, even if you check in the replica. Item properties on a replicated document are always overwritten by the Replication Daemon when it updates the replica.

Note The Title value of a replica document is supposed to be blank if it is blank in the source document. But if the following conditions apply, you will instead see the target Item ID of the replica document in place of a blank Title value:

- Replication Services is shut down on the source library before you add the document.
- You add the document to the source library using IDM Desktop, assigning it a Doc Class containing an optional Title parameter that you leave blank.
- Replication Services is restarted on the source library.

Replication Home Server

Administrative domains offer centralized management of objects from the Domain Master library. However, if the original folder or document object resides on an in-domain member library, you cannot perform all the management tasks from the Domain Master that has a replica of the object. Tasks, such as deletion of the object or granting access rights to the object, can be performed only from the library where the object originated or was sourced.


To overcome this problem and ensure centralized management of documents and folders from the Domain Master, the **Replication Home Server** feature is used on the Domain Master. On using this feature, the Domain Master is referred to as **Home Server**.

Centralized management of documents and folders

When the Replication Home Server feature is implemented on one or more replication folders on the Domain Master, the folder is referred to as a Home Server enabled replication folder. This feature ensures that all the documents and subfolders within a replication folder are automatically sourced to Home Server. In other words, all management tasks related to these subfolders and documents are possible only from Home Server.

Any subfolder or document that you add to a Home Server enabled replication folder will be sourced on Home Server, regardless of the library system where it was originally added. If a document or a folder is added from a member library to a Home Server enabled replication folder, it is first replicated only to Home Server. The original document or folder in the member library system is then converted to a replica and Home Server stores the original document or the folder. The replication of this document or folder to the remaining target libraries is then initiated from Home Server.

The following figure illustrates how documents and subfolders are managed when the Replication Home Server feature is enabled on a replication folder.



Note In the above figure, reference to the term "object" is restricted to documents and subfolders.

Adding subfolders and documents on an offline Domain Master

If the Domain Master is offline due to reasons, such as network problem, Replication Home Server feature allows you to add objects to Home Server enabled replication folders, on the Domain Master, from any member library participating in replication. When the Domain Master is restored online, the newly added document or folder is automatically sourced to the Domain Master and replicated to all target libraries from the Domain Master.

Applicable only to member libraries within the administrative domain

The Replication Home Server feature is applicable only within an administrative domain. You can source only those documents and subfolders to Home Server that you add from the member libraries and not from outside domain libraries. However, an outside domain library can act as a target library to replicate the documents or subfolders from Home Server

Refer to “Enabling Replication Home Server” on page 625 for instructions on enabling the Replication Home Server feature on replication folders.

Setting Up Replication Services

Replication Services is automatically installed. However, to use Replication Services, you must enable and configure the service.

To set up, configure, maintain, and use Replication Services, you need the following:

- Two or more participating library systems with unique names. Contact FileNet Professional Services if your site has library systems with the same name.
- The participating library systems must use the same database type and release version, MSSQL or Oracle. FileNet does not support replicating to a different database type.
- The participating library systems must use the same library system version. If your site requires a software patch, the same patch must be installed on all participating libraries.
- For administrators, install the Content Services Admin Tools on at least one administrative client.
- For users, install a client interface to the library systems, such as IDM Desktop.

To set up Replication Services

1. Plan your replication requirements carefully. Refer to [“Planning for Replication Services” on page 601](#).
2. Review your replication plan with FileNet Technical Support. The support technicians can help you avoid potential configuration problems. Refer to Support Programs after logging in to the [FileNet Worldwide Customer Support](#) web site.
3. Back up all library systems that will participate in Replication Services. Be sure to archive the resulting backup files accordingly. (For details, refer to [“Restoring Content Services” on page 339](#).) This precaution ensures that you can recover fully from configuration mistakes such as setting up incorrect administrative domains.
4. Ensure that Replication Services is stopped on all participating library systems.
5. Configure the Oracle database for Replication Services. Refer to [“Editing the tnsnames.ora File for Replication Services” on page 621](#).
6. (Recommended) Copy/create a common set of custom properties, Document Classes, and CVLs on the participating library systems. The custom properties must match in name, data type, length, parent object, and multivalue setting. The CVLs should match in datatype.

If custom properties do not match, those properties are not replicated. If the duplicate CVLs do not match on datatype, the values are still merged, and unpredictable results occur. For example, if a numeric CVL is merged with a string CVL of the same name, the string values will produce errors on the system defined to use numeric values.

- To copy custom properties, use the COPY_CP utility on the System Administration Tools CD to copy custom properties. Refer to [“COPY_CP” on page 382](#) for more information.
- To copy Document Classes and CVLs:

Use the DocClass/CVL Tool export feature. For more information, refer to the Content Services Admin Tools documentation.

7. Enable Replication Services in either of two ways:
 - On every server containing the Replication Services component of a library system that will participate in Replication Services, select the Enable Replication check box on the Replication tab of the CS Configuration Tool.
 - Using CS Admin Tools, log on to each library system that will participate in Replication Services and set the Replication Enabled? property to Yes in the System object of the library system.
8. (Optional) Establish the administrative domain. Refer to [“Administrative Domains” on page 606](#). Add all appropriate member libraries.
9. Configure Replication Services settings. You can specify log file settings, activity intervals, folder and document title settings, and default access rights for replicas.
 - For Windows, use the CS Configuration Tool. Refer to [“Replication Services Configuration Parameters” on page 624](#) and the *CS Configuration Tool Online Help* for more information.
 - For UNIX, edit the `$IDMDS_HOME/dsconf.ini` file. Refer to [“Replication Services Configuration Parameters” on page 624](#) for more information.
10. (Optional) If you are using an administrative domain, start the Library Services on the Domain Master. Then start each member library system one at a time. After adding each library, wait one or two replication cycles. Then review the replication log (**repl.log**) and take appropriate action if needed. Add additional libraries one at a time as needed.
11. Configure the root-level replication folders. Refer to [“Configuring Replication Folders” on page 623](#).

Allowing for Replication Latency

When a user changes an item in a library, a work order is generated in the system to duplicate the change in the specified replication libraries. For example, when a user drags a document into a replication folder, a work order is generated to add a copy of the document to the specified replication libraries. Network traffic, library system down time, and Replication configuration parameters affect the delay between the time a change is requested and when the work order completes processing on all systems.

There is no direct feedback from the system that alerts the user or the administrator that the work order has been completed. The only way to determine when the changes have been fully propagated across all participating replica libraries is to log into each library and confirm that the new or updated item exists in each library.

When the user makes additional changes to a document or item, more work orders are created and processed. If a user attempts to make additional changes to a document or item before the earlier work order is completed, the new work orders can become stuck in the work order queue and block all processing in the queue.

For example, if a user adds a document to a replication folder, and then immediately attempts to move or delete the document before the document is replicated to all libraries, the replication queue may hang. The system attempts to begin processing the second work order, but since the document does not yet exist in all libraries, the move or delete operation cannot proceed. The user can still see and manipulate the original source item, but the replication process halts until FileNet Technical Support can intervene.

If you are using an administrative domain, the latency period also applies to changes to users and groups.

Replication latency also effects indexing on the replica libraries. Users will not see accurate searches on documents until after the indexes on the target libraries have been updated.

Wait for changes to replicate across all library systems before requesting additional changes.

Configuring Latency

The Content Services Configuration Tool Replication Tab settings for Wakeup Seconds, Remote Interval, and Run Times affect the latency time between when a change is requested and when the change appears on all participating libraries.

For best results, use the default configuration settings. When replication runs continually all day, there is usually enough time to complete all the work orders in the replication queue.

If you change the configuration settings, make sure that you allow ample time for all replication work orders to complete. If the window for Run Time is too small, then all the work orders in the queue may not be processed, and your libraries will never be synchronized.

For very low-volume environments (less than 20 items updated per day), you can limit the Run Time to off-peak hours.

Bulk loading documents, folders, or user changes (in an Administrative Domain) can significantly slow down your system. If you need to bulk load thousands of documents for replication or make a large number of changes, you should wait until off-peak hours.

A document can arrive at a replica library before its folder does. You can search for the document in the library even if the folder has not yet been created.

Configuration Results

Replication Services configuration information is stored in the following locations:

| Windows | UNIX |
|--|--|
| <i>drive:\filenet\repl</i> | <code>\$IDMDS_HOME/librarysystem/repl/pbin</code> |
| <i>drive:\filenet\repl\logs\librarysystem\repl.log</i> | <code>\$IDMDS_HOME/librarysystem/repl/pbin/repl.log</code> |
| Windows Registry | <code>\$IDMDS_HOME/dsconf.ini</code> |

Replication Services also does the following:

- (For Windows installations) Creates subdirectories for each library system participating in Replication Services. Initially, there will be only a subdirectory for the library system specified for this installation. When the administrative client specifies the additional libraries participating in Replication Services, the program creates subdirectories for each of these library systems.
- Creates a ReplDaemon user in the library system. The ReplDaemon user has Admin access rights to the library system, which must be retained so that Replication Services performs as expected. Do **not** modify the ReplDaemon user password.

Editing the tnsnames.ora File for Replication Services

For Replication Services to work with an Oracle database, you need to modify the **tnsnames.ora** file.

Look in the **\orant\network\admin\sqlnet.ora** file and check whether the following line is present:

```
names.default_domain =value
```

If it is not, do not make any changes. If the names.default_domain line exists, open the **\orant\network\admin\tnsnames.ora** file with a text editor and append the server name(s) with *.value* (where, by default, the *.value* entry is *.world*).

For example, you would change the following lines:

```
librserve =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS =
      (PROTOCOL = TCP)
      (HOST = librserve)
      (PORT = 10040)
    )
  )
  (CONNECT_DATA = (SID = saro)
)
```

to read instead as:

```
librserve.world =
(DESCRIPTION =
  (ADDRESS_LIST =
    (ADDRESS =
      (PROTOCOL = TCP)
      (HOST = librserve)
      (PORT = 10040)
    )
  )
  (CONNECT_DATA = (SID = saro)
)
```

where *librserve* is the host name of the property server. In your environment, you would replace *.world* with the domain name associated with the particular Oracle database used by the library system.

Establishing a Replication Administrative Domain

Caution Plan your domains carefully before you configure Replication Services. You cannot undo your changes to an administrative domain.

- A library system can belong to only one administrative domain.

- Once assigned to an administrative domain, a library system cannot be removed from that domain. Refer to “[Planning for Replication Services](#)” on page 601 and “[Administrative Domains](#)” on page 606 for more information.
- The users and groups assigned to any library system in an administrative domain must have network access to all library systems in the domain.
- Verify that you performed the setup steps outlined in “[Setting Up Replication Services](#)” on page 618.

To establish a replication administrative domain:

Note For tool-specific instructions, refer to *CS Explorer Online Help*.

1. Back up all library systems that will be in the administrative domain.
2. Make a list of the library systems that will be in the domain, and complete the Administrative Domains worksheet. Refer to “[Administrative Domains Worksheet](#)” on page 657.
3. Verify that the Storage Manager and Content Search Manager for the participating library systems are running.
4. Decide which of the participating library systems to designate as the Domain Master library, and start Replication Services for this library only.
5. Start CS Explorer and log in to the Domain Master library as a member of the Administrators group.
6. Select Tools > Replication Domain Administration, and add libraries to the domain. You must have Admin access rights to any library system that you include in an administrative domain.

Caution When adding a library system to a replication (administrative) domain via the Replication Domain Management tool, do not modify the CS-generated Comment field, which appears in the Other Systems list entry for the added library system. In particular, if you modify the Comment field by removing the bracket-delimited identifier, the added library system will not appear in the Domain Members list. That is, if you exit the Replication Domain Management tool and then launch it again, the added library system will not be in the list, even though it acts as a normal library system in the domain.

7. Select one of the following password modes for this administrative domain to specify how users in this administrative domain can change their passwords and how (or whether) these changes are synchronized across the member library systems:

- Source Library

Users can change their password only when logged in to their source library system. When a user changes a password, Replication Services replicates it across all other member library systems.

Note There will be some delay in effecting this change across participating library systems. This will depend primarily on the synchronization settings for the replication service at each library system.

- Any Library

Users can change their password when logged in to any member library system. When a user changes a password, Replication Services replicates it across all other member library systems.

Note Problems can occur if conflicting password changes are made at more than one member library system within the latency period necessary to replicate the change.

- Library Specific

Users can create distinct passwords for each member library system. Users can change their password for any member library system and have it apply only to that library system. When you add a new member library system to this administrative domain, the passwords for all users replicated to the new member library system are set to their source library system password assignment.

8. Select one or more available libraries to add to your domain. If the library you need does not appear in the list, click Add System to define the missing library.
9. Add the selected library to the list of Domain Members.
10. Click Done to confirm the domain membership assignments.

Configuring Replication Folders

The term Replication Folder refers to a root-level folder that the administrator uses to designate the participating libraries where documents and subfolders are replicated. The Replication Daemon copies this folder and its contents to all the participating library systems. Users with appropriate access rights can add documents and subfolders to the replication folder from any participating library.

IDM Desktop displays the replication folders and subfolders with this icon:



You can name the replication folder in any manner you choose. You can include information that identifies the folder as a replication folder, or you can use the replication configuration setting for Folder Title Format to add information to the replicated folder title. For more information, refer to “[Replication Services Configuration Parameters](#)” on page 624.

You can designate multiple replication folders, each referencing different target library systems. By designating multiple replication folders and library system associations, you can determine the distribution of replicas. For example, you might have one replication folder that references library systems A, B, and C, another replication folder that references library systems C and D, and another that references library systems B and D.

You can create a replication folder even if Replication Services is not enabled for that library (that is, the Replication Enabled? property in the System object is set to No). If users add documents to such a replication folder, the documents are queued for replication. Once you set the Replication Enabled? property to Yes and the Replication Daemon is started, the documents are replicated.

Note For tool-specific instructions, refer to *CS Explorer Online Help*.

To create a replication folder:

1. Verify that the appropriate library systems participating in Replication Services are running and accessible.
2. Start CS Explorer and log in as a member of the Administrators group.
3. Select Tools > Folder Administration, and at the root (top) level add a new folder or select an existing folder to designate as a replication folder.

If you choose to use an existing folder, all the existing subfolders and documents will be replicated. This may lead to unintended consequences. Refer to [“Managing Replication Folders”](#) on page 629 for more information.

4. Right-click this folder, and then select Replication.

A dialog box opens, listing the library systems participating in Replication Services. If the library you need does not appear in the list, click Add System to define the missing library.

5. Add the selected library to the Replicate To list. If needed, you can remove a library system from the Replicate To list.

Note If you want to enable Replication Home Server on the replication folder, select the Enable Home Server check box. This will ensure that all documents and subfolders within the replication folder are sourced on the Domain Master. To know more about this feature, refer to [“Enabling Replication Home Server”](#) on page 625 and [“Enabling Replication Home Server”](#) on page 625.

6. Click Done.

Replication Services automatically creates a copy of the replication folder on each of the specified library systems, and the Replication Daemons replicate any subfolders or documents added to the replication folder.

Note There will be some delay in effecting the replication across participating library systems. The delay depends primarily on the synchronization settings for Replication Services on each library system. Do not rename or delete a replication folder while it is being replicated to participating library systems. For more information, refer to [“Replication Services Configuration Parameters”](#) on page 624.

Replication Services Configuration Parameters

You can configure these replication settings for a library system:

- Interval and frequency of Replication Daemon activity
- Log file name, location, size, and related options
- Information to include in titles of replicated folders and items
- Storage and archive categories
- Index settings
- Default access rights for replicas
- Default online limit

For a library system running on Windows, the Replication Services configuration settings are saved in the Windows Registry. You use the CS Configuration Tool to control replication configuration options for a library system. Refer to [“Replication Tab”](#) on page 169 for more information.

Enabling Replication Home Server

Replication Home Server can be enabled on one or more replication folders on the Domain Master. When you enable Replication Home Server on a replication folder, the entire replication tree, including the subfolders and documents within the replication folder, will inherit the characteristics of this feature.

Refer to “[Replication Home Server](#)” on page 616 to get an overview of the Replication Home Server feature.

Notes

- You can enable Replication Home Server only on replication folders existing on the Domain Master. If you try to enable Replication Home Server on a member library system, you'll get an error message.
- Only Administrators can enable Replication Home Server on a replication folder.

There are two ways to enable Replication Home Server

- You can enable Home Server using Folder Administration window while configuring the existing folders for replication.
- You can enable Home Server on already replicating folders using the HMSRVMIG command utility.

Using the Folder Administration window

The Folder Administration window can be used in any of the following situations:

- A root folder exists on the Domain Master and is not replicating to any library. You want this folder to replicate to one or more member libraries and want that any new document or subfolder added to this replicating root folder should be sourced on the Domain Master.
- A replication folder on the Domain Master is replicating only to outside domain libraries. You want to include one or more member libraries in the replication and want that any new documents and subfolders added to the replication folder be sourced on the Domain Master.

Before you start

- Identify the replication folder(s) on the Domain Master on which you want to enable Replication Home Server.
- Identify the target libraries that you want to include in replication.
- Make sure that you include at least one member library in the list of target libraries.

To enable Replication Home Server using the Folder Administration window

1. Start CS Explorer and log in as an administrator to the Domain Master library where you want to set up Home Server enabled replicated folder. From the Tools menu, select **Folder Administration...** The Folder Administration dialog box opens.
2. Select the root folder on which you want to enable Replication Home Server and select **Replicate** from the Context menu.

The **Folder Replication Settings** dialog box is displayed.

3. Select the **Enable Home Server** check box.

Note This check box is disabled if you try to enable Replication Home Server on a root folder existing on a member library or an outside domain library. The check box is also disabled for replication folders on the Domain Master that are already replicated to member libraries.

4. A list of library systems is displayed. From this list, select the library systems, where you want the selected folder to be replicated. If you do not find the required library in the list, click the Add System button to add the required library.
5. Click OK.

Using the HMSRVMIG utility

The HMSRVMIG utility enables you to source all the documents and subfolders within a replication folder from member libraries to Home Server. This utility is useful when the source of the documents and subfolders within a replication folder is scattered across member libraries and the Domain Master.

The HMSRVMIG utility is used in any of the following situations:

- A replicated folder exists on a member library and this folder is already getting replicated to other libraries in the domain. You want this folder, its levels and sublevels, and any new additions to this folder to be sourced on the Domain Master.
- A replicated folder exists on the Domain Master and is already replicated to other member libraries in the domain. You want this folder, its levels and sublevels, and any new additions to this folder to be sourced on the Domain Master.

Pre-requisites for running the HMSRVMIG utility

Before you run the utility, you should ensure that the following pre-requisites are fulfilled:

- All target libraries should be available and connected to Home Server: All the documents and subfolders within a replication folder have to be sourced on Home Server. All the target libraries for this replication folder must to be updated about the swapping of status between Home Server and the original member library. It is necessary that all target libraries are available and connected to Home Server, else the sourcing of documents and subfolders to Home Server will not be completed.
- User logons should be disabled on all the participating libraries: Disabling user logons ensure that users do not perform any operations on any document or folder marked for migration.
- No pending work requests for any document or folder that has to be migrated: Any pending work request for a document or folder to be migrated, will fail if the document or folder is migrated before the request is processed. This will cause inconsistent object state between the member library and Home Server.

Note You can use the ReplMon utility to check for any pending work requests.

- Create backups for all target libraries: You should have backups of all target libraries participating in the replication of the replication folder on which you plan to run the utility. In the absence of backups, you will not be able to restore the libraries to the original state, if the utility does not complete successfully.
- (Recommended) Run the Home Server Check (HMSRVCHK) utility: This utility checks all the documents and subfolders within the replication folder and generates a list of all subfolders and documents that will be sourced on Home Server on running the HMSRVMIG utility and the documents and subfolders that have exceptions.

The subfolders and documents with exceptions will not be sourced on Home Server. Once the list is available, you should resolve the exceptions for the documents and subfolders that you want to source to Home Server. (Refer to “HMSRVCHK” on page 420 to learn more about this utility and how to run this utility.)

Note If you resolve the exceptions after running the HMSRVMIG utility, the subfolders and documents with exceptions will not be sourced on Home Server, till you run the HMSRVMIG utility again.

Working of the HMSRVMIG utility

While running the HMSRVMIG utility, you must specify the ID of the replication folder on which you want to enable Home Server. This is irrespective of the folder existing on a member library or Home Server, you need to specify the ID assigned to that folder on Home Server.

The utility checks all the subfolders and documents within the replication folder, to find their source library and perform one of the following:

- If an object is sourced on the Domain Master, the object is left unaffected.
- If the utility finds an object sourced on a member library, it selects that object for sourcing it to Home Server. The object status is swapped between the member library and Home Server. After the swapping, all target libraries are updated for the change in the source of the object.
- If the utility finds an object sourced on an outside domain library, it considers this object as an exception to migration. Such object is not sourced on Home Server; instead the original object is left on the outside domain library with the replica on the Domain Master.

The following figure depicts how the utility handles different documents and subfolders.

Managing Replication Services

Typical Replication Services management includes these tasks:

- Stopping, starting, and disabling Replication Services. Refer to “Starting, Stopping, and Disabling Replication Services” on page 636.
- Managing users and groups. Refer to “Managing Users and Groups” on page 633.
- Managing Document Classes and CVLs. Refer to “Managing Document Classes and CVLs” on page 600
- Configuring Replication Folders. Refer to “Configuring Replication Folders” on page 623.
- Managing replication folders. Refer to “Managing Replication Folders” on page 629.
- Modifying document or folder replicas. Refer to “Modifying Document or Folder Replicas” on page 632.
- Replicating documents. Refer to “Managing Documents” on page 634.

Managing Replication Folders

Replication Services replicates documents and folders that are placed into a replication folder. The system administrator designates the root-level replication folder, and users with Author access rights can add documents and subfolders to it.

Typical folder management includes these tasks:

- Replicating subfolders
- Moving a replication folder
- Removing a replication folder

Caution After placing a folder into a replication folder, do not move, rename, or delete the folder until it and all of its contents have been fully replicated.

Replicating Folders

Users and group members can drag existing folders to the replication folder or create new subfolders. The user or group must have with Author access rights to the replication folder. When a user drags an existing folder into a replication folder, all of the original folder's subfolders and documents are moved as well.

Proceed with caution when moving existing folders. In particular, you should bear in mind that:

- A moved folder hierarchy may contain subfolders or documents to which the owner of the parent folder has access rights of None. The user would have no way of knowing that such subfolders or documents exist in the hierarchy. Replication Services will replicate these subfolders and documents to the other library systems. Depending upon the replication configuration settings, these subfolders and documents may be visible in IDM Desktop.
- If you accept the default item access list configuration for Replication Services, all documents in the folder hierarchy are replicated and given the default item access list of the ReplDaemon user (no more than Viewer access rights for all users by default).

- If you ever delete the root-level replication folder or modify the libraries being replicated to, the subfolders and documents will not be visible in IDM Desktop. The documents can be located with a search, but you will need to create a new directory structure. For more information, see “Removing a Replication Folder” later in this section.

Replicating Folder Structures

For best results, make all changes to your folder structures on the original library. While users can add subfolders from a replica library, this practice will cause problems if you need to de-replicate the libraries. If your replication system becomes unstable, FileNet Technical Support may request that you de-replicate libraries for troubleshooting.

- Any subfolders added to the folder structure from a replica library will be deleted from the original library on de-replication. You will need to manually recreate the lost folder structure on the original library. Additionally, the root folder only exists on the original library, and when you de-replicate, the folder structure on the replica library is lost. The subfolders that were added from the replica library are lost on the replica library, and you will have to recreate either a partial or full folder structure on the replica library. Any documents that originated on the replica library are placed into the library’s root folder, and users will not be able to see them using IDM Desktop. You can search for the lost documents and recreate the folder structure, but this process can be time consuming.

Access Rights for Replicated Folders

The users and groups in the access list of a replicated folder are limited to maximum access rights of Author. The ReplDaemon user of the participating library system always has Owner access rights.

Inside and outside an administrative domain, Replication Services sets the user and group entries for the folder replica's access list depending on how you configure the Replication Services settings. In general, you can enter a setting that makes the entries in the folder replica's access list match one of the following:

- The access list of the original folder, plus the ReplDaemon user.
- General users are Viewers, plus the ReplDaemon user.

For more information, refer to “[Replication Services Configuration Parameters](#)” on page 624.

On a Home Server enabled replication folder, users can perform the following operations:

| Domain Master | Member library | Outside domain library |
|----------------------------------|--|------------------------|
| Member library operations, plus: | Outside domain library operations, plus: | Add |
| Delete | Check out | View |
| Modify access rights | Check in | Print |
| File | | Copy |
| Unfile | | |

For more information on a Home Server enabled replication folder, refer to “[Replication Home Server](#)” on page 616.

Moving a Replication Folder

A user or a member of a group having Owner access rights to a subfolder can move the original subfolder within the Replication folder tree, but not out of the Replication folder tree. That is, you can move original subfolders up or down in the hierarchy, but the root-level folder must remain the same. If you try to move a replicated subfolder, you will receive an error message.

You cannot move a replica subfolder or document. Also a subfolder containing replica subfolder or replica document underneath its tree hierarchy cannot be moved.

You can determine if a folder is a replica or an original from its ID. An original replication folder or subfolder has a numeric ID. A replicated folder's ID begins with R. For example, the Planning folder in the Engineering library has the ID 943389629, and the Planning folder in the Marketing library has the ID R38440d9a. In this case, the folder in the Engineering library is the original folder.

Note For Home Server enabled replication folders, the folder ID does not denote whether folder is a replica or an original. Refer to [“Replication Home Server” on page 616](#) for more information.

You can determine the library systems for a replication folder by viewing the root-level folder properties in IDM Desktop. The Comments field lists the library systems that a folder is replicated to. The first library system listed is the original library system for the root-level folder. Subfolders do not indicate replication libraries in the Comments field. Remember that subfolders will replicate to the same libraries as the root-level replication folder.

Note You cannot move an existing root-level replication folder into another folder.

Removing a Replication Folder

You cannot remove a replicated subfolder directly from its target library. Instead, you can delete the original subfolder or you can dereplicate the root-level replication folder.

When you delete the original subfolder from its library system, the replica subfolders and replica documents are deleted from all participating library systems.

When you dereplicate the root-level replication folder, you remove the association between the folder and the other library systems. You can remove the association for a single library system or for all participating library systems. These rules apply to the replicated documents and folders on the system that no longer participates:

- Any subfolder or document that originates from other library systems is removed.
- Any document that originates from the removed system remains on its original system. Users can access the document only by running a search. The document replicas are removed from the other systems and are no longer accessible.
- All subfolders that originate from the removed system are removed from all library systems.

Notes

- There will be some delay in effecting the removal across participating library systems. The delay depends primarily on the synchronization settings for the replication service at each library system. For more information, refer to [“Replication Services Configuration Parameters” on page 624](#).
- Inside an administrative domain, users can check out document replicas independently of the source document. If a document replica is checked out at the time the replication folder is removed, the checkout operation is automatically cancelled.

- You can delete documents or subfolders within a Home Server enabled replication folder only from the Domain Master and not from any library system. To know more about a Home Server enabled replication folder, refer to “[Replication Home Server](#)” on page 616.
- For tool-specific instructions on how to dereplicate a replication folder, refer to *CS Explorer Online Help*.

Modifying Document or Folder Replicas

A replica can be modified under certain conditions, based on whether or not it originates inside an administrative domain and whether the user has at least the necessary Author access rights.

Inside Administrative Domains

Users can modify document or folder replicas if all of the following conditions are true:

- The replica's original object exists on a library system within the administrative domain.
- The user exists on both library systems inside the administrative domain and has adequate access rights (Author) to modify the replica.
- The user checks out the replica document before making the changes and checks in the replica after making changes.

For document replicas, you can make changes to the version properties. You must check out and check in the replicated document to replicate your latest version property changes back to the original document and other replicas. If you do not check in your changes to the version properties, they are never replicated, and they will be overwritten when the Replication Daemon updates the document when someone else checks in changes.

Any changes you make to item properties on a replica are temporary, even if you check in the replica. Item properties on a replicated document are always overwritten by the Replication Daemon when it updates the replica.

Outside Administrative Domains

Outside an administrative domain (or if no administrative domains exist) the only users who can modify replicas are those currently working in the Administrators group in the associated library system. Any changes an administrator makes are temporary.

There are restrictions to the kinds of modifications that an Administrator can make to the replica:

- All changes are temporary and will be overwritten by the Replication Daemon when it updates the replica.
- You cannot delete a replica.
- You cannot change the access list to provide access rights greater than the maximum allowed (Viewer) for any user or group.
- You can make changes to a replica in a target library, such as renaming a document. However, the changes are not replicated to the other library systems, and the changes will be overwritten the next time a user updates the original document.
- If you change a replica's object properties, the changes are not replicated to the other library systems. For example, you could add library-specific users or groups to the access list for a particular document or folder replica.

Managing Users and Groups

Users and groups are only replicated inside an administrative domain. Outside of an administrative domain, you must manually add users and groups to all participating library systems.

Inside an administrative domain, when you add new instances of User, Group, Group Member, Default Item Access List, and Access Control objects on any library system, Replication Services automatically replicates those objects across all participating library systems. The Administrators group and group member entries for the Administrators group are not replicated.

When a user or group of the same name exists on more than one library system inside an administrative domain, Replication Services will designate the user or group with the earliest creation date as the original object and the remaining user or group objects as replicas.

For example, if a user named Danika exists on two library systems inside an administrative domain, the Danika user with the earliest creation date is designated the original User object instance, and the other Danika user is designated a replica of this object instance.

Caution Be sure to check that user and group names existing in different library systems actually belong to the same individuals (or sets of individuals, in the group case). Otherwise, the Access List and the Default Item Access List will assign unanticipated access rights during replication to particular user or group replicas. For example, Replication Services could inadvertently assign user or group replicas Author rights to documents and folders they should not be able to control or even view. Use the user/group utility USGRPCHK to check for duplicate user and group names across libraries.

For best results, use the Domain Master to add new users and groups to the administrative domain. If possible, use the Domain Master to manage users and groups. If the Domain Master library system or the initial storage server on which it resides is inoperable, Replication Services will not be able to replicate the changed users or groups to the other library systems in the administrative domain until the Domain Master is available.

Inside an administrative domain, the Access Control object instance for a user or group is also replicated, with maximum access rights set to Owner.

Caution If you add a user and a group to a domain master, allow them to replicate to all target libraries, and then add the user to the group, be sure to verify that the group memberships have been replicated to the target libraries before deleting the group on the domain master. This means you should wait at least as long as the value of Remote Interval (see “The CS Configuration Tool” on page 164), after adding the user to the group, before deleting the group. But note that some group memberships may not replicate to all target libraries within the time specified as the Remote Interval. If you delete the group before the group memberships have been replicated, the Member Add objects (one for each target library, corresponding to adding the user to the group) will remain indefinitely in the replication queue. If this happens, contact [FileNet Worldwide Customer Support](#) for assistance in restoring the integrity of the replication queue.

Modifying User and Group Replicas

Inside an administrative domain, you can make temporary modifications to some properties, such as Default Storage Category or Default Archive Category on a user or group replica. However, the property values you change on a replica are not reflected in other library systems - neither in the original User or Group object instance, nor in the User or Group object replicas existing in other library systems inside the administrative domain. The changes will be overwritten the next time the Replication Daemon updates the replica.

Note While you can set the “Add Users?” and/or “Add Groups?” properties to Yes for a source user object, the changes to these properties will not be propagated to replicas of the user object. For security reasons, the properties will always be set to No for user object replicas.

If you want changes to User and Group object instances to be replicated inside an administrative domain, you must make your changes to the original object. For best results, use one library system to add new users and groups to the administrative domain. Use the Domain Master to manage users and groups.

Any property value you change on the original user or group object will cause the Replication Daemon to automatically replicate the updates to all User or Group object replicas. Any previously modified properties on the replica object are overwritten with the associated property value from the original object.

Deleting User and Group Replicas

Once replicated to the library systems inside the administrative domain, the User and Group object replicas will continue to exist on the library systems until you specifically remove them.

Note For tool-specific instructions on how to modify a user or group replica, refer to *CS Explorer Online Help*.

Managing Documents

Users with Author access to a Replication folder can replicate documents. When the user drags a document to the Replication folder (or its child folders), the Replication Daemon automatically creates the document replicas on the target library systems.

Users can add and move documents to the target Replication folder just as they can to any other folder. However, if a user adds a new document directly to a Replication folder only and you later remove the Replication folder, the document is no longer visible in a folder. You can locate the document with a search. Refer to “[Managing Replication Folders](#)” on page 629 for more information.

In case you have enabled Home Server on a replication folder, you can add document(s) to this replication folder from any target library. However, if the document is added from a member library participating in replication, then the source of the added document will always be on Home Server, that is, the Domain Master. The member library from which it was originally added will have the replica of the document. Refer to “[Replication Home Server](#)” on page 616 for more information.

Access Control Lists

The access list for document replicas is determined by the Replication Services configuration settings. Replication Services handles access control lists (ACLs) differently inside and outside of an administrative domain.

| Outside an Administrative Domain | Inside an Administrative Domain |
|--|--|
| <p>The access list for the document replicas matches either:</p> <ul style="list-style-type: none"> The entries from the default item access list of the ReplDaemon user on the target library system (with maximum access rights of Viewer), plus the ReplDaemon user from the target library system with Owner access rights. -Or- The access list of the original document (with maximum access rights of Viewer), plus the ReplDaemon user from the target library system with Owner access rights. <p>Replication Services ignores a user/group entry if it does not exist on the target library system. Otherwise, the Replication Daemon copies all other recognized user/group entries to the access list of the document replica on the target library system.</p> | <p>The access list for the document replicas matches either:</p> <ul style="list-style-type: none"> The entries from the default item access list of the ReplDaemon user on the target library system (with maximum access rights of Author), plus the ReplDaemon user from the target library system with Owner access rights. -Or- The access list of the original document (with maximum access rights of Author), plus the ReplDaemon user from the target library system with Owner access rights. <p>Replication Services copies a user/group entry if it does not exist on the target library system. Within the domain, users and groups are replicated to all library systems.</p> |

For more information on replication access list configuration settings, refer to [“Replication Services Configuration Parameters” on page 624](#).

Moving an Original Document

When a user moves or copies an original document to another folder in the Replication tree, the document replicas on the target library systems are also moved or copied to the corresponding folder. You must have Author access to the replication folder to move the original document.

Removing an Original Document

Only users with at least Owner access rights to the original document can remove the associated document replicas. Once a user removes the original document, Replication Services removes all document replicas on the participating library systems as well.

Note There will be some delay in effecting the removal across participating library systems. The delay depends primarily on the synchronization settings for the replication service at each library system. For more information, refer to [“Replication Services Configuration Parameters” on page 624](#).

Checkout and Checkin Operations with Document Replicas

Inside an administrative domain, a user with Author access to the document replica can check out that replica if the original document is not already checked out by another user.

For users with IDM Desktop, it is possible to make an unconfirmed check out. The user does not know that the check out is tentative, but the replication log will contain the message “Tentative checkout cancelled.” When the user attempts to check in the document, the check in option is not available.

This tentative check out occurs only in these conditions:

- When the user attempts to immediately check in the document, within about 5 seconds of the check out.
- When the Replication Daemon is not running on either the original library or the library where the replica resides at the time the user requests the check out.
- When there is a network interruption and the Replication Daemon cannot access the library where the original document resides.

To avoid tentative checkout scenarios, a user can check out the document on the original library. Checkouts on the original library are always guaranteed.

Outside the administrative domain, these operations require access to the original document on its original library system:

- Check out
- Check in
- Cancel checkout

Note You can make changes to item and version properties for original documents. However, you must check out and check in the original document from the original library to replicate your changes to the respective target libraries. If you do not check in your changes, they are never replicated to the target libraries.

For document replicas inside an administrative domain, you can make changes to the version properties when you check in the document. If you do not check in your changes, they are never replicated, and they will be overwritten when the Replication Daemon updates the document.

Any changes an administrator makes to item properties on a replica are temporary. Item properties on a replicated document are always overwritten by the Replication Daemon when it updates the document.

Starting, Stopping, and Disabling Replication Services

You can stop and restart Replication Services to perform maintenance tasks, or you can completely disable Replication Services for a library System.

The `dsstrt`, `dsstop`, and `dsinfo` utilities enable you to control a Replication Services instance by starting, stopping, and providing information on Replication Services. Library systems running in Windows can use the CS Service Manager or the Services utility in the Control Panel. For more information, see [“Utilities” on page 369](#).

To disable Replication Services for a library system:

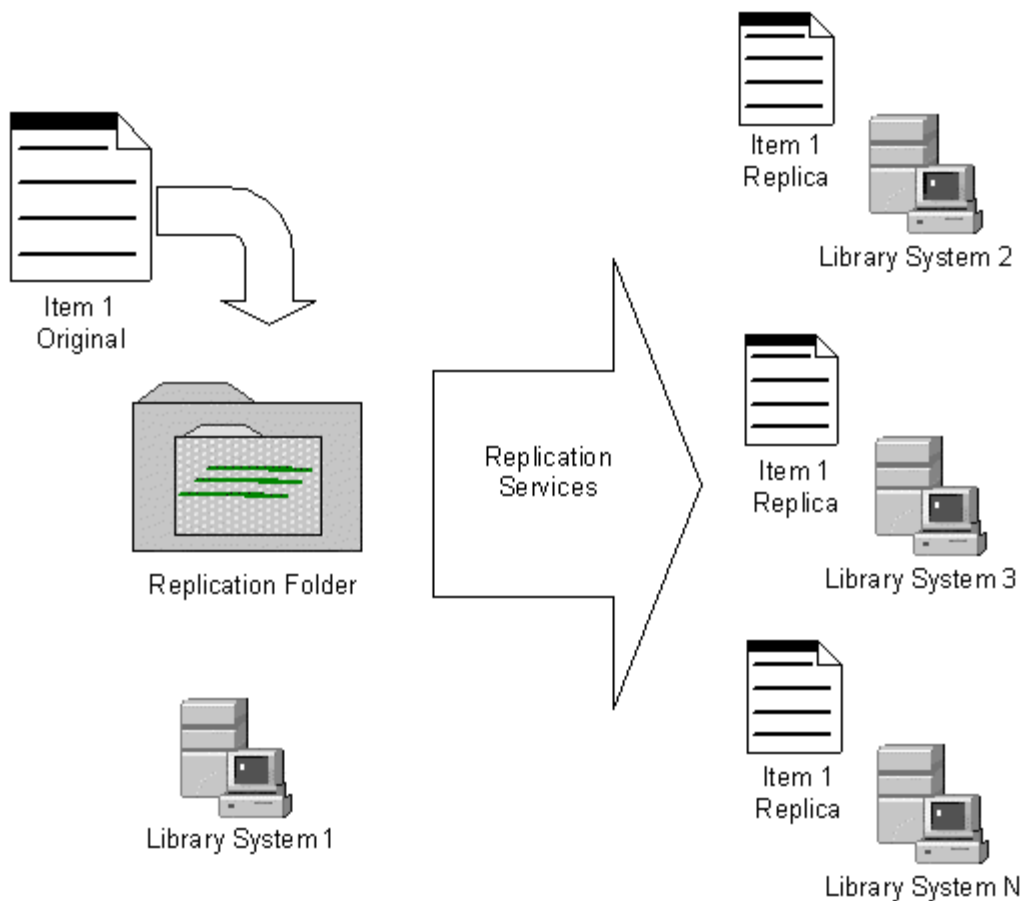
1. Notify all users to log out of the library system.
2. Use CS Explorer to log in to the library system as an Administrator.
3. Disable logons for this library system. This ensures that users cannot move or modify replicas before you disable Replication Services.
4. In the System object, set the Replication Enabled? property to No.
5. Stop Replication Services by using `dsstop` or the CS Service Manager.
6. Enable logons for this library again so that users can access the library.

Replication Guidelines for End Users

As a user, you need to know if a document is to be replicated and whether replicas of it appear in other libraries. These guidelines are intended to help you understand the impact of replication on your documents and folders.

Your system administrator can create special folders in a library called *replication folders*. The subfolders and documents that you and others add to a replication folder are automatically copied to other library systems. Your administrator determines which folders are replication root-level folders and which libraries each root-level folder is copied to. The subfolders and documents are always replicated to the same libraries as the root-level replication folder.

Note There can be one or more replication folders for a library system, with each replication folder referencing different library systems.



Replication Services Overview

IDM Desktop displays all replication folders and subfolders with this icon: 

Warning The replication process takes time to complete copying documents, folders, and information between libraries. Changes you make may not appear in IDM Desktop for a while. After you add a document or subfolder to a replication folder, **wait** for the change to appear in the replica libraries before you make any other changes.

For example, if you add a document to a folder, **do not** attempt to delete or move the document until it appears in all replica libraries.

Replication Terms

The following terms are used to describe the concepts and behaviors in Replication Services.

| Term | Meaning |
|-----------------------------------|---|
| Original Object | The original object that is replicated to other libraries. Modifications made to a original object are replicated to all replica objects automatically by the Replication Daemon. Original objects were called source objects in previous versions of IDMDS. |
| Original Library | The library system where the original object is added. Original documents and original folders are created or added on the original library system. Original libraries were called source libraries in previous versions of IDMDS. |
| Replica Object Target Object | The replicated (copied) object. Also referred to as a replica. For example, a folder replica or document replica. Modifications made to a target object are not always replicated back to the original object or any other replicated objects by the Replication Daemon. |
| Replica Library Target Library | The library system where replicas are located. The replica is created and owned by the Replication Daemon on the target library. |
| Replication Daemon | The replication software responsible for maintaining consistency between original objects and target objects. |
| Administrative Domain | <p>A collection of library systems that share security features. Within an administrative domain, any changes to User, Group, or Group Member objects are automatically replicated to the library systems. Stored searches created with the System Administration Tool, Document Classes, and CVLs are also replicated across the domain.</p> <p>Administrative domains are useful in environments where original documents may be located in a remote location and modifications are performed locally.</p> <p>Your site administrator will tell you if your site uses Administrative Domains.</p> |
| Domain Master | The library system where an administrative domain is created and maintained by the system administrator. |
| Outside an Administrative Domain | All library systems that are not members of an administrative domain and that are participating in replication. Users can view, print, or copy replica documents. Users make all modifications on the original document (in the original library system), and all other libraries contain view-only document replicas. |

| | |
|---------------------------------|---|
| Inside an Administrative Domain | Replication features that apply only within an administrative domain. Users can check out, check in, and modify any document replica (whose original document is inside the administrative domain) from any library system within the domain, provided that the user has the appropriate access rights. |
| Replication Folder | Refers to a root-level folder that the administrator uses to designate where documents and subfolders are replicated. Users can add documents and subfolders to the replication folder. For best results, always add documents and subfolders from the original library. In IDM Desktop, participating libraries are listed in the folder's Comment field. |

What information is copied between libraries?

The information copied between libraries depends on how your system administrator has configured your site. Your system administrator determines which folders are copied to which libraries, and determines if your site uses administrative domains to add additional features.

The following information is copied between all libraries that participate in Replication Services:

Always copied between libraries

- Folders
- Documents
- Compound Documents
- Annotations for a Document
- Published Documents
- Versions
- Custom property values when the same property exists on both libraries

If your administrator has configured an Administrative Domain for replication, then the following items are copied between the libraries inside the domain (in addition to standard items):

Only copied inside an Administrative Domain

- Users
- Groups
- Group Memberships
- Searches stored created with the System Administration Tools
- Default Item Access List
- Document Classes when those properties are identical on both libraries

- CVLs (Controlled Vocabulary Lists)


To determine if your site uses Administrative Domains:

- Ask your administrator.

Warning The replication process takes time to complete copying documents, folders, and information between libraries (inside or outside of an Administrative Domain). Changes you make may not appear in IDM Desktop for a while. After you add a document or subfolder to a replication folder, **wait** for the change to appear in the replica libraries before you make any other changes.

Identifying replica folders and libraries

IDM Desktop displays replication information for replication folders and documents. To view folder or document properties, right-click the folder or document, and then click **Properties**.




Folder Properties

In the example to the left, the replication folder named Proposals exists on the library named Saloon and replicates to the library named Zipper on the server named TESTBERT.

The ID tells you if the folder is an original or a replica. The ID of an original folder begins with a digit; the ID of a replica folder begins with an 'R'. For example, if the ID is 94294465, then the folder is an original.

The **Comment** field lists all the libraries that are configured for replicating the root-folder.

The **Replication Libraries** field lists the replica libraries and servers for this folder.



Document Properties

In the example to the left, the document named JanuaryReport exists on the library named Saloon as a replicated document. It is replicated from the library named Boomtown on the server named DAWSON.

The **Library** field lists the library where this document is located.

The **Replicated From** field lists the original document's Item ID and the original library name and server.

Adding subfolders and documents to a replication folder

Your system administrator must create the top-level replication folder (root folder) and specify which libraries the folder replicates to. Once the root folder exists, you can add folders and documents if you have the necessary permissions to access the folder. For best results, always add folders and documents from the original library system of the root-level replication folder.

All of the objects (subfolders and documents) placed in a replication folder are replicated to the other library systems. For folders, this may be a single folder or a hierarchy (tree) of folders.

Note You must have Author access to the replication folder to add documents or folders for replication. Your system administrator configures your access level.

To add a subfolder to a replication folder:

1. Log on to the original library system for the root-level replication folder.

2. Click the desired folder and drag it into a replication folder. You can add the folder to the root-level folder or to a subfolder. A replication folder has this icon:



Warning The replication process takes time to complete copying documents, folders, and information between libraries (inside or outside of an Administrative Domain). Changes you make may not appear in IDM Desktop for a while. After you add a document or subfolder to a replication folder, **wait** for the change to appear in the replica libraries before you make any other changes.

To add a document to a replication folder:

1. Log on to the original library system for the root-level replication folder.
2. Click the desired document and drag it into a replication folder. You can add the document to the root-level folder or to a subfolder. A replication folder has this icon:



Warning The replication process takes time to complete copying documents, folders, and information between libraries (inside or outside of an Administrative Domain). Changes you make may not appear in IDM Desktop for a while. After you add a document or subfolder to a replication folder, **wait** for the change to appear in the replica libraries before you make any other changes.

Removing subfolders and documents from a replication folder

When you remove a subfolder from a replication folder, all of its documents and subfolders are removed from the replica library systems. The original documents always remain in their original library system.

Notes

- While it is not recommended, you can add subfolders to a replication folder from a replica library. If you add subfolders from a replica library, then later remove the subfolder from the replication tree, the folder structure on the replica libraries and the original library are both affected. Folders added from a replica library are deleted from both the replica library and the original library for the root-level replication folder. The documents remain in their original library, but you cannot see them in IDM desktop. You can still search on these documents, but you will need to recreate the folder structure.
- Do not try to drag or copy a Replication subfolder outside of its root Replication folder. You can only move a replication subfolder within the same root-level Replication folder structure. That is, you can move a Replication subfolder into another Replication subfolder in the same folder tree, but you cannot move it to another folder tree that has a different root folder.
- The root-level replication folder can only be deleted by your system administrator.

To remove a subfolder from a replication folder:

1. Log on to the original library system for the replication folder. You must have Author access to the replication folder to remove subfolders from replication.
2. Right-click the desired folder, and select **Delete** from the context-menu.
3. In the confirmation dialog box, click **Yes**.

Warning The replication process takes time to complete copying documents, folders, and information between libraries (inside or outside of an Administrative Domain). Changes you make may not appear in IDM Desktop for a while. After you add or remove a document or subfolder to a replication folder, **wait** for the change to appear in the replica libraries before you make any other changes.

To remove a document from a replication folder:

1. Log on to the original library system for the original document. You must have Author access to the replication folder to remove documents from replication.
2. Click the desired document and drag it from the replication folder into a regular folder. The document will be removed from the replica libraries after a short delay. It remains in the original library.

Warning The replication process takes time to complete copying documents, folders, and information between libraries (inside or outside of an Administrative Domain). Changes you make may not appear in IDM Desktop for a while. After you add or remove a document or subfolder to a replication folder, **wait** for the change to appear in the replica libraries before you make any other changes.

To delete a document from all libraries:

1. Log on to the original library system for the original document. You must have Author access to the replication folder to remove documents from replication.
2. Right-click the desired document, and select **Delete** from the context-menu.
3. In the confirmation message box, click Yes. The document is deleted immediately from the original library, and after a short delay, it is deleted from all the replica libraries.

Working inside an administrative domain

If your site uses administrative domains, you can checkout, modify, and checkin replica documents and original documents from any library system in the administrative domain. Your site administrator will tell you if your site uses administrative domains.

Working with compound documents in replication folders

You can replicate compound documents across libraries. For best results, follow these guidelines:

- Keep the parent document and all of its child documents in the same folder in the same original library.

Note In case the child documents for a compound document are not kept in the replicated folder, the links between parent and child documents will not be maintained in the replica.

- Make all changes on the original library (even inside an Administrative Domain).

It is safe to check out and modify the content of a parent or child document from any library where it appears, but changes to the linkage between documents (including OLE embedding) should only be made on the original library. Since some types of links may not stand out upon casual inspection, it is safest to make all changes on the original library.

Managing Replication Home Server

The operations that can be performed on a Home Server enabled replication folder depends on the library system from which the operations are performed and the access rights available to the user.

Operations allowed on a Home Server enabled replication folder

The following operations are allowed for users on a Home Server enabled replication folder:

| Domain Master | Member library | Outside domain library |
|----------------------------------|--|------------------------|
| Member library operations, plus: | Outside domain library operations, plus: | Add |
| Delete | | View |
| Modify access rights | Check out | Print |
| File | Check in | Copy |
| Unfile | | |

Disabling Replication Home Server

You can disable Replication Home Server on a replication folder by deselecting the Enable Home Server check box in the Folder Replication Settings dialog box. After disabling Replication Home Server, the existing documents and subfolders will continue to have their source on Home Server. However, any new document or subfolder added to the replication folder will not be sourced on Home Server. The source of the newly added document or subfolder will be the library on which it is added. Also, the replication of all new documents and subfolders will initiate from the library where these are added and not from Home Server.

Note Only Administrators can disable Replication Home Server on a replication folder from the Domain Master.

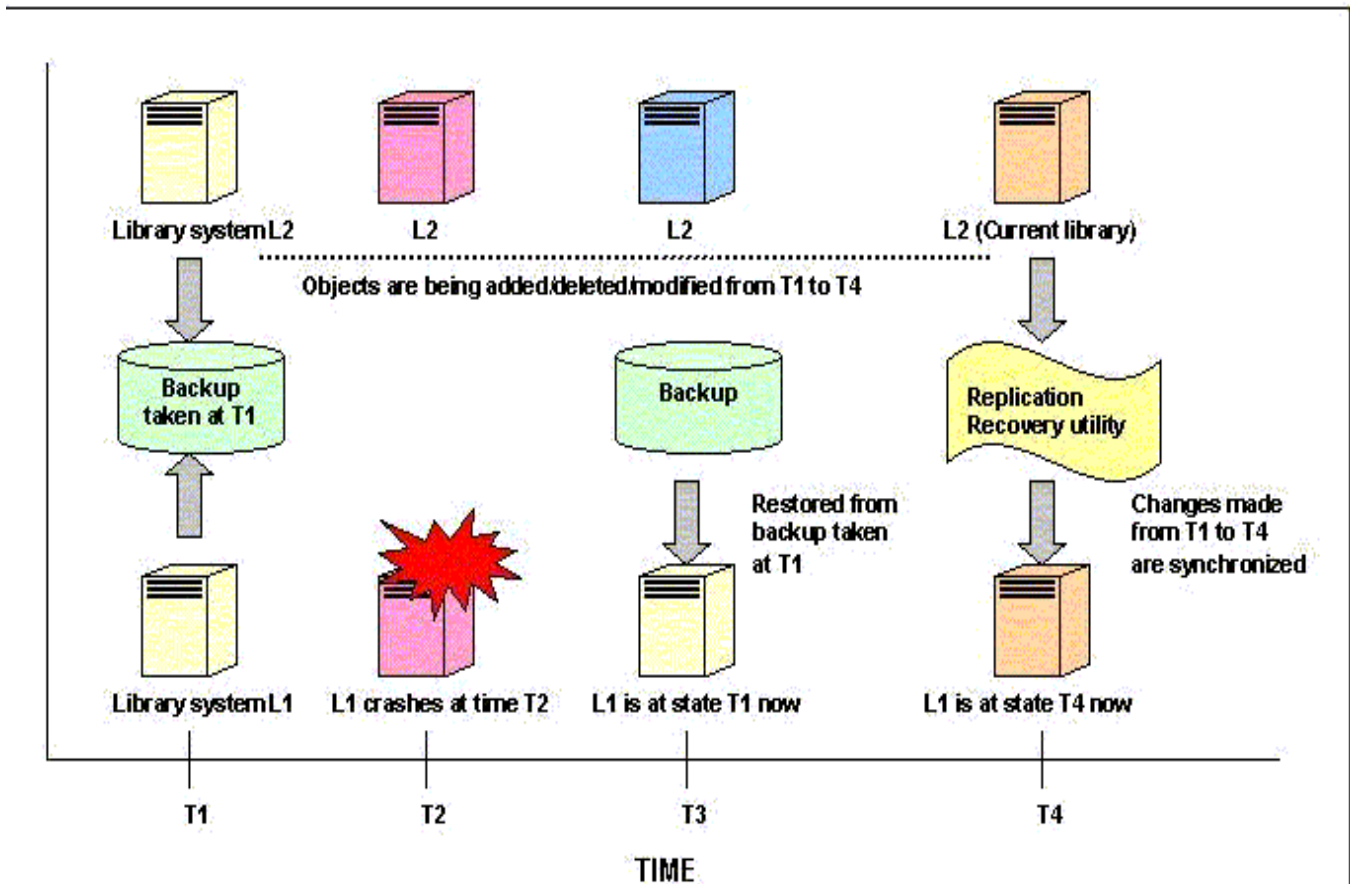
Replication Recovery

The goal of any recovery process is to ensure that all committed data prior to a system failure is fully recovered and reflected once the system has been restarted. In situations such as failure of a library system's host server, it might be necessary to restore all library systems in the administrative domain from their respective backups to maintain inter-library data integrity. However, this will result in loss of data created/modified since the last backup. Replication Recovery utility overcomes this situation and helps you recover the affected library system. Once you have restored the affected library system from backup, this utility rolls-forward the library state to synchronize the replicating data with other libraries in the administrative domain.

The Replication Recovery utility compares the replicating objects on the restored library system with a specified library system in administrative replication domain to identify the discrepancies. It then fixes these discrepancies to recover the restored library system and synchronize it with the current state of other member libraries participating in the administrative replication domain.

Note While running the utility, the library from which you want to recover the crashed library is referred to as the **Current library**. The damaged library system that needs to be synchronized is referred to as the **Recovered library**.

The following illustration demonstrates the use of Replication Recovery utility:



The Replication Recovery utility performs synchronization of all replication objects such as, folders, documents, users, groups, versions, and object relationships. For more information on objects that would be synchronized, refer to [“Objects Supported for Replication” on page 597](#)

Working of the Replication Recovery Utility

- If the replica of an object is found missing from the Recovered library, then the object is re-replicated so as to create a replica on the Recovered library.
- If an object other than folder or document is available on the Recovered library and has been deleted from the Current library, it is deleted from the Recovered library as well.
- If a source folder object is available on the Recovered library and missing from the Current library, then the utility de-replicates it from the Recovered library and moves it along with its contents to a folder named *RR Lost+Found*. The utility creates this folder for a specific date on the Recovered library to store folders found to be missing on running the utility on that particular date.
- If a document is available on the Recovered library and missing from the Current library, then the utility unfiles it from all the replicating folders. Documents sourced on the Recovered library continue to exist but are not filed in any folder. Replica documents that are unfiled from all replicating folders are deleted from the Recovered library.
- If any of the properties of a replicating object are found to be different on Current and Recovered libraries, then the properties are updated on the Recovered library.
- On running the utility, if a user object is found missing on the Recovered library, the utility adds that user object on the Recovered library and sets the password for the user to that of the Current library.
- There are certain configurable properties for the document and folder replica objects, that can be controlled using CS Configuration tool e.g. Item title format, Replicate Item ACLs etc. The comparison and updates for the folder & document properties between the original and replica object are performed based upon the latest configurations.
- All temporary version checkouts on the Recovered library are performed to a folder named "RRCHECKOUT". This folder is created at the location specified for the **Fix Report**. Refer to [“Fix Report” on page 647](#) for more information. On completion of the Replication Recovery utility, the working files for all documents checked out on the recovered library are available in this folder.

Note Only one instance of the Replication recovery utility can be run at a time.

Reports Generated

Synchronization Report

The Replication Recovery utility generates a Synchronization report that contains a list of all discrepancies detected on the Recovered library while comparing replicating objects with the Current library. The Synchronization report is a preview file showing the transactions that would be performed during fixing the discrepancies. A unique record ID identifies each discrepancy in the report. A sample row in the Synchronization report is displayed:

FileNet Content Services Replication Recovery Report.
Started: 02/05/04 16:52:50

Invoked by
User name: Admin
System: winsqlib03
Server: fn-cs2w2k

Current Library
User name: Admin
System: winsqlib01
▶ Server: fn-cs2w2k

Recovered Library
User name: Admin
System: winsqlib03
Server: fn-cs2w2k

02/05/04 16:52:53 -- [1] Folders, with folder id 1075979175 on current library winsqlib01 and folder id T85160008 on recovered library winsqlib03, exist with differing properties. These properties differ: Parent Folder, HomeServer Status, DateTime1. The differing properties will be made identical. The parent id T85160004 for the folder on recovered library winsqlib03 also differs and will be updated.

02/05/04 16:52:53 -- [2] Folders, with folder id 1075979176 on current library winsqlib01 and folder id T85160009 on recovered library winsqlib03, exist with differing properties. These properties differ: Parent Folder, HomeServer Status, DateTime1. The differing properties will be made identical. The parent id T85160008 for the folder on recovered library winsqlib03 also differs and will be updated.

Each row in the Synchronization report contains:

- Date and time when the discrepancy is reported.
- Unique ID provided to the discrepancy, and
- Discrepancy information and the action to be taken to fix the discrepancy

Fix Report

Replication Recovery utility generates a Fix report that contains the details and status of the synchronization operations performed for removing the discrepancies detected in the Synchronization report. The record IDs used for discrepancies in the Fix report are same as the record IDs used in the Synchronization report. However, the order of occurrence of discrepancies in the two reports may be different. A sample row in the Fix report is displayed:

► FileNet Content Services Replication Recovery Fix Report.
Started: 02/05/04 16:53:02

Invoked by
User name: Admin
System: winsqlib03
Server: fn-cs2w2k

Current Library
User name: Admin
System: winsqlib01
Server: fn-cs2w2k

Recovered Library
User name: Admin
System: winsqlib03
Server: fn-cs2w2k

```
02/05/04 16:53:03 -- [1] Update of properties and parent id for folder T85160008 on
recovered library winsqlib03 was successful.
02/05/04 16:53:03 -- [2] Update of properties and parent id for folder T85160009 on
recovered library winsqlib03 was successful.
```

Each row in the Fix report contains:

- Date and time when the discrepancy is reported.
- Unique ID provided to the discrepancy in the Synchronization report, and
- Discrepancy information on whether the action taken to fix the discrepancy has succeeded or failed.

Note The Synchronization reports and Fix reports are generated at the specified location under their corresponding name of *syncreport.log* or *fixreport.log*. If a file by that name currently exists in the specified directory, a backup of the existing file is created by appending the current date and time information in the *yymmddhh[AM/PM]MMSS* format followed by the *.bak* extension. For example, if the Synchronization report named *syncreport.log* already exists in the specified path, the utility renames it as *syncreport03102410AM2500.bak*.

Error Logs

In addition to Synchronization and Fix reports, the Replication Recovery utility also generates the following error logs:

- Synchronization Error Log (syncreport.err)
- Fix Error Log (fixreport.err)

The error logs provide a quick reference to determine whether any errors were encountered during the Replication Recovery discrepancy detection or fixing process. These error logs are stored at the same location where the Synchronization and Fix reports are stored.

Preparing for Replication Recovery

1. Backup the library systems.

Perform a complete backup of all library systems in the administrative replication domain except the recovered library system (refer to “[Backing Up the Library System](#)” on page 336).

2. Restore the recovered library system.
 - a. Restore the last available backup of the Recovered library. Refer to “[Restoring Content Services](#)” on page 339 for more information.
 - b. Put the Recovered library system in the single-user mode using CS Admin Tools. A library system is in single-user mode when
 - The Administrators group is your active group.
 - The Logons Enabled? property in the System object is set to No.
 - No other session but yours is open.

WARNING Do not start the Storage manager, Content Search and Replication services for the recovered library at this time.

3. Identify the Current Library.

Replication objects other than documents and folders can be synchronized by choosing any library in the administrative replication domain as the Current Library, as these objects are replicated to all libraries in the domain. Documents and Folders can be replicated between specific libraries in the administrative domain instead of all participant libraries.

4. Process pending work orders on Current library.
 - Using the Replication tab of the Content Services Configuration tool, Set the **Run Times** to the default settings and **Remote Interval** to 1 for Current library. Refer to “[Configuring Content Services](#)” on page 164 for more information
 - Stop and start Replication Services for the Current library. Also ensure that the Storage Manager service for the Current library has been started.
 - Ensure that there are no pending work requests for the Current library in the Administrative domain. In case there are any pending work requests, the utility will wait for these requests to be completed.

Note You can use the ReplMon utility to check for any pending work requests.

5. Synchronize custom properties

Run the COPY_CP utility to copy the custom properties from the Current library to the Recovered library.

COPY_CP utility is available on the System Administration Tools CD. Refer to “[COPY_CP](#)” on page 382 for more information.

6. Start services on Recovered library

Start the Storage Manager and Content Search services for the Recovered library. Do not start the Replication service.

Running the Replication Recovery Utility

1. Start CS Explorer and log in to the Current library, with the credentials of the default administrative user. From the Tools menu, select **Replication Recovery**. The Replication Recovery dialog box opens.

Note Alternatively, you can double-click on the CSReplRecovery.exe located in the directory where the CS Administration tools have been installed.

2. Click the Select button displayed with Current to open the Current Library dialog box. In this dialog box, you specify the name of the Current library system and the name of the server on which the Current library exists.
3. Click OK to close the Current Library dialog box.
4. Specify the username and password of the default Administrative user on the Current library and click OK.
5. Repeat steps 2 to 4 for specifying the Recovered Library.
6. Specify the directory path where the replication recovery utility will generate the Synch Report containing the discrepancies that needs to be fixed. By default, the utility stores the report in the directory where the CS Administration tools are installed.
7. Specify the directory path where the replication recovery utility will generate the Fix Report containing the discrepancies fixed by the utility. By default, the directory path is where the CS Administration tools are installed.

Note You must have the Write access on the path(s) that you specify for generating reports.

8. Select the Logging option for Synchronization and Fix reports. The logging option determines the level of details to be included in reports.
 - 0: Errors only ensures logging of the discrepancy/fix details along with the errors encountered.
 - 1: (SPI Details) ensures logging of the invoked CS API routines in addition to the events logged in the option '0'.
 - 2: (SQL Details) ensures logging of the database SQL activities in addition to the events logged in the option level '1'.
9. Click Generate Synch report to generate the discrepancies report. You can skip this step if you want to directly fix the discrepancies without viewing the report containing discrepancies.
10. Click Fix Discrepancies to fix the discrepancies detected on the Recovered library system.

While the utility is generating the Synchronization report, the count of objects processed is displayed in the status bar of the Replication Recovery dialog box. During the fixing process, the status bar displays the progress on the count of discrepancies fixed.

The utility also reports the errors encountered while detecting and fixing discrepancies in error log files. Refer to [“Error Logs” on page 648](#) for more information.

Note Once you have recovered the library system by running the Replication Recovery utility with all the required member libraries as the Current library in succession, you should start the Replication service for the Recovered library. This will ensure complete synchronization of objects that have been added, deleted, or changed on other libraries while the utility was running.

Replication Recovery Utility Limitations

The following are the constraints of the Replication Recovery utility:

- Once you start running the utility, you cannot pause, cancel, or abort the utility.
- Original documents created by the replication recovery utility have ACLs similar to replica: For a replica object, the maximum access rights, any user other than ReplDaemon can get is of an "Author". As a result the information regarding the level from which the ACLs have been downgraded is not available on the library having the replica object. Therefore, the original objects created by the replication recovery utility have downgraded level of ACLs as available on the Current library. Administrators need to specifically upgrade the ACLs on the Recovered library for the created original objects, if required.
- The utility does not synchronize the security (ACLs/DIALs and memberships) and properties for system-defined groups and users. You need to synchronize these manually on the Recovered library, if required.

Replication Planning Worksheets

Use the replication planning worksheets to help you gather and analyze the existing data before you configure Replication Services. The worksheets can help you identify duplicate objects and avoid overwriting important information. The worksheets also provide a baseline for future changes.

If you intend to use administrative domains, these worksheets will assist you in planning a successful replicated environment. Remember, once a library is added to an administrative domain, it cannot be removed from the domain.

Refer to the [“Configuration Worksheets” on page 43](#) that you used to plan your library systems for information on your existing systems.

After completing the worksheets, contact FileNet Technical Support to review your plans before you implement any changes. Technical support can help ensure a smooth and successful transition to using Replication Services.

| Worksheet | Description |
|--|---|
| “Replication Users and Groups Worksheets” on page 653 | Lists the users and groups from each library system that will participate in Replication Services. |
| “Replication Document Class and CVL Worksheet” on page 655 | Lists the Document Class and CVLs from each library system that will participate in Replication Services. |
| “Administrative Domains Worksheet” on page 657 | Lists the library systems that belong to an administrative domain and the Domain Master. Provides a check list for preparing to configure an administrative domain. |

Replication Users and Groups Worksheets

Use this table to help you identify potential duplicate users and groups. List all the participating library systems across the top of the table. Then list each participating user name and group name in the first column. Place a check mark in the library column where the user or group account exists. Use the user/group utility USGRPCHK to detect duplicate users and groups across libraries.

Locate the users and groups that have more than one check mark. Ensure that the duplicate names on different library systems refer to the same individual or group members. If not, rename one of the duplicates before configuring Replication Services. Otherwise, the wrong person may access inappropriate documents and cause a potential security breach.

You may discover that the same individual has more than one user account. Rename all accounts to the same name before configuring Replication Services.

| User & Group Names | Library 1 | Library 2 | Library 3 | Library 4 |
|--------------------|-----------|-----------|-----------|-----------|
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

| User & Group Names | Library 1 | Library 2 | Library 3 | Library 4 |
|--------------------|-----------|-----------|-----------|-----------|
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

Replication Document Class and CVL Worksheet

Use this table to help you identify potential duplicate Document Classes and CVLs.

1. List all the participating library systems across the top of the table. Then list the Document Classes and CVLs in the first column. Place a check mark in the library column where the Document Class or CVL exists. Use the Document Class and CVL utility DTCVLCHK to detect duplicate Document Classes and CVLs.
2. Locate the Document Classes and CVLs that have more than one check mark. Ensure that the duplicate Document Classes and CVLs on different library systems are identical. This includes standard and custom properties in the Document Class, whether the property is required, default value (if any), and assigned CVLs. When reviewing CVLs, ensure that the data type and CVL data values are identical.
3. If not, either modify both entries so that they become identical, or create a new Document Class and perform a bulk modify to the existing documents for that library and indicate the new Document Class name before configuring Replication Services.

For ease of management, create all DocClasses and CVLs on the Domain Master. Use the DocClass/CVL Tool export feature to create export files for importing on the Domain Master. For more information, refer to the Content Services Admin Tools documentation.

Remember, inside an administrative domain, Document Classes and CVLs are automatically replicated. Outside an administrative domain, Document Classes and CVLs are not replicated.

| DocClass and CVL Names | Library 1 | Library 2 | Library 3 | Library 4 |
|------------------------|-----------|-----------|-----------|-----------|
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

| DocClass and CVL Names | Library 1 | Library 2 | Library 3 | Library 4 |
|-----------------------------------|------------------|------------------|------------------|------------------|
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |
| | | | | |

Rendition Services



Welcome to FileNet's Rendition Services for Portable Document Format (PDF). It is an optional program that creates PDF renditions of documents stored in your CS library system. You can use Rendition Services to render documents from many different file formats in PDF. Once the documents are in PDF, other users can read them in Acrobat viewers or Web browsers. The document and its PDF rendition are each stored as separate items in the library system.

Rendition Services uses Lquent's rendition engine to create PDF renditions of documents stored in your document repositories. Rendition Services uses special plug-ins and the document authoring programs to produce an accurate rendition of the document. Rendition Services handles all the processing, including the communication with your FileNet property server.

Note Rendition Services and the publishing functionality in IDM Desktop support general compound documents but are not designed to fully support automatic maintenance of OLE links in those documents. As a workaround, when you want to make changes to a child document portion of an OLE compound document (for example, to an Excel spreadsheet linked into a Word document) you must always check out and open both the parent and the affected child document to refresh the link. Only then should you check the parent and child documents back into the Content Services library. You can then republish the document to PDF format.

Rendition Services Components

Rendition Services contains the following components:

- LiquePDF Printer is a printer driver that uses the Windows Print function to produce a print stream, which encapsulates page elements in PDF.
- Plug-ins are format translators that let you open source documents in authoring programs and print them in PDF.
- CS 5.5.0 and IDM Desktop lets you access repositories. It handles retrieving and distributing documents and their renditions, and gathering properties for documents.
- Business Services are designed to perform and manage the following processes:
 - retrieving documents, renditions and properties
 - rendering to PDF
 - distributing to a CS library
- Vista Administration is a status viewer that lets you monitor job details and status as well as administer the Business Services running on that server.
- RenderPerfect Administrator is a utility that allows you to change queueing times and days or configure Rendition Services to render files on local drives. Job tickets are also defined using this utility. Rendition Services controls document processing using job tickets. A job ticket is an XML file that defines the documents to be processed and the processes to be performed.

Hardware and Software Requirements for Rendition Services

To verify that all hardware and software requirements are satisfied on all servers where you will install Rendition Services, refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#).

Rendition Services can convert documents from many different authoring programs and formats. Before you can create renditions, you must install a copy of the authoring program on the server. Refer to [CS 5.5.0 Windows Oracle Platform Requirements](#) at [FileNet Worldwide Customer Support](#) for a list of the supported authoring programs.

In addition to installing the authoring programs, you must also make sure the authoring programs are installed as the service account user. Refer to [“Prepare a Server”](#) on page 662 for more information on service accounts.

Keep the following points in mind:

- You need to install only the authoring programs for the documents that you will be processing using Rendition Services. Be sure to install and test the programs for all formats you intend to convert before installing Rendition Services.
- Follow the installation instructions carefully for each authoring program. Rendition Services needs the programs to run error-free in order for it to convert and publish documents accurately.

Uninstall Previous Versions of Rendition Services

CS 5.5.0 is being released with Rendition Services 5.5.0, which is an upgrade and replacement for kPublisher 1.1.x. Only upgrades from kPublisher 1.1.x are supported. To upgrade Rendition Services that was based on kPublisher 1.1.x, do the following:

1. Uninstall Acrobat Server, Adobe Plug-ins, and any previous versions of Rendition Services. Refer to the latest CS 5.4 documentation on the [FileNet Worldwide Support](#) web site. Refer to “[About the Documentation](#)” on [page 12](#) for instructions on how to access updated Content Services documentation.
2. Reboot your system.
3. Prepare your server and install Rendition Services 5.5.0. See “[Install Rendition Services on Servers](#)” on [page 664](#) for details.

Installation Overview

Use this checklist to guide you as you install and use Rendition Services.

| Step | Reference |
|--|---|
| Prepare your server. | This task includes setting up a service account and installing Adobe Acrobat Distiller. Refer to “ Prepare a Server ” on page 662 for more information. |
| Install Rendition Services on servers. | This task includes doing the actual Rendition Services install. In addition, information for configuring the Rendition Services server, domain and applications is also provided. Refer to “ Install Rendition Services on Servers ” on page 664 for more information. |
| Complete the server setup. | This task includes setting your startup properties for the server, setting your VistaTempDir permissions, modifying the Lipient PDF printer, configuring the Business Service Manager, and setting custom access properties. Refer to “ Complete the Server Setup ” on page 671 for more information. |
| Disable options in Microsoft Word and Corel WordPerfect. | This task includes information on disabling options in Microsoft Word and Corel WordPerfect while using Rendition Services. Refer to “ Prepare Microsoft Word Documents ” on page 677 and “ Prepare Corel WordPerfect Documents ” on page 678 for more information. |

| Step | Reference |
|--|---|
| Configure the Corel WordPerfect printer. | This task includes information on configuring the Corel WordPerfect printer. Refer to “Configure the Corel WordPerfect Printer” on page 678 for more information. |

The following checklist provides optional procedures you can do with Rendition Services.

| Optional Steps | Reference |
|---|---|
| Configure server content types. | This tasks includes information for associating plug-ins with content types. Refer to “Configure Server Content Types” on page 680 for more information. |
| Check the status of Rendition jobs. | This task includes information on using Vista Administration. Refer to “Check the Status of Rendition Services Jobs” on page 703 for more information. |
| Control Rendition Services jobs. | This task includes information on using RenderPerfect. Refer to “Control Rendition Services Jobs” on page 706 for more information. |
| Specify archived jobs. | This task describes the process of dedicating a machine to perform a job archiving service. Refer to “Specify Archived Jobs” on page 714 for more information. |
| Use watched folders and apply IDM Desktop Publishing preferences. | This task includes information on using watched folders with Rendition Services and the Desktop Publishing preferences applied on the Rendition Services server. Refer to “Watched Folders and IDM Desktop 4.0 (or higher) Preferences” on page 714 for more information. |
| Work with log files. | This task includes information on managing the logs Rendition Services generates while processing jobs. Refer to “Log Files” on page 715 for more information. |
| Change the RenderPerfect login. | This task includes information on setting up login information if you plan on using RenderPerfect. Refer to “Change the RenderPerfect Login” on page 716 for more information. |
| Uninstall Rendition Services on a server. | This task includes information on uninstalling Rendition Services on a server. Refer to “Uninstall Rendition Services on a Server” on page 717 for more information. |

Prepare a Server

Go through the following list before installing Rendition Services on a server or multiple servers:

1. Make sure the server(s) meets the minimum hardware and software requirements. Refer to “[Hardware and Software Requirements for Rendition Services](#)” on page 659 for a list of requirements.
2. Create a service account for installing all applications, including authoring programs, and Rendition Services 5.5.0. The service account you use to install Rendition Services 5.5.0 must be a member of the Windows 2000 domain, and must also be a member of the local Administrators group on that server. The user must also have a password. Check to see that the server has all the necessary network connections and privileges.

In a production environment, a single service account should be created to install all applications and install, configure and run Rendition Services. The account can be a domain account or a local account if you intend to install Rendition Services on a single server.

If you are installing Rendition Services on a single server, you need to create only a local account.

If you are installing Rendition Services on multiple servers, the service account must be a member of your Windows 2000 domain.

Note Use the same service account to install and configure multiple Rendition Services on all domain computers.

3. Install and test all authoring programs you intend to use with Rendition Services. Use the same service account on all domain computers that will render and enrich documents. Refer to “[Hardware and Software Requirements for Rendition Services](#)” on page 659 for a list of supported authoring programs.

Make sure to install authoring programs as “all users” if you are given a choice.

4. Exit and close any other programs running on the server, including virus-protection programs.
5. Make sure your Microsoft SQL Server 2000 or Oracle 8i is set up properly. Test your connection to your SQL or Oracle database. Do an ODBC test for SQL and tnsping for Oracle. In addition, be sure a valid ODBC DSN is set up using the Oracle ODBC driver and not the Microsoft Oracle ODBC driver.

Note Although Access database is listed as an option, using an Access database is not supported in this release.

The Oracle database to which you want to connect must already exist; you cannot create it in the Vista Configuration application as you can a SQL database.

Note You must run a script as the database system administrator before configuring the Oracle database. The Oracle-repository.txt script file is located in the root of the install directory. These scripts create the schema, objects, table space and users required for Rendition Services as well as populating the schema.

For more details about creating an Oracle database, see your Oracle documentation.

To run the required scripts

- a. Locate Oracle-repository.txt in the root level of the Rendition Services install directory.

Note You must be the system database administrator before you can run the scripts.

- b. Use SQLPlus or another Oracle database tool to run the Oracle-repository.txt script.

6. Make a list of your library systems that will be publishing documents with Rendition Services. Verify that you can access all these CS libraries using IDM Desktop installed on the Rendition Services server.

Note It is required you assign the library name to be the same name assigned for the system name when configuring a CS library for IDM Desktop. Do not use any aliases.

Install Adobe Acrobat Distiller Server 5.0.5

Rendition Services uses Adobe Acrobat software to render documents in PDF format.

To install the Adobe Acrobat software

1. Insert your Rendition Services CD into your CD drive.
2. Locate and double-click the **Adobe Acrobat Distiller Server 5.05** folder.
3. Double-click the appropriate language directory, such as English, and then double-click the **Installer** folder.
4. Double-click **Setup.exe** to start the install.
5. In the Adobe Acrobat Distiller 5.0.5 Setup dialog box, click **Next**.
6. Choose the language in which you want to use with the software, and then click **Next**.
7. In the Software License Agreement dialog box, click **Accept**.
8. Select **Typical** for your setup preference, and then click **Next**.
9. In the User Information dialog box, enter your information, and then click **Next**.
10. Enter the following serial number:
WAW500R7107001-900-287
11. Click **Yes** to confirm your registration information is correct.
12. Review your current settings, and then click **Next**.
13. Click **OK** once the installation is complete.

Install Rendition Services on Servers

Rendition Services can be installed on a single server or multiple servers. If you are installing and configuring Rendition Services to run on multiple servers, you should perform a normal install on each server and specify the same database for all Rendition Services servers.

After Rendition Services is installed, you also need to configure the Rendition Services server using Vista Configuration.

Notes

- Collocating Rendition Services with a CS server is not supported.
- Using terminal services during an installation is not supported.
- To avoid publishing failures, the user installing Rendition Services must be logged in at all times on the Rendition Services server. It is okay for the desktop to be locked, but the user must be logged in.

Rendition Services Install

Note If you are installing Rendition Services on multiple servers, refer to “[Multi-Server Install](#)” on page 675 for additional information.

To install Rendition Services

1. Log in to the domain as the service account user you created while preparing your server.
2. Close all open programs including virus checking.
3. Insert the Renditions Services CD into the computer's CD-ROM drive.

Note If the autorun does not start, click **Start > Run**, type **d:\setup.exe**, where d is the letter of the drive containing the Rendition Services CD, and then click **OK**.

4. In the Welcome dialog box, click **Next**.
5. In the Customer Information dialog box, enter your user name and organization information.
6. Specify the location in which you want to install Rendition Services, and then click **Next**. A default location is displayed. Click **Browse** if you want to change the location.
7. Review the current settings before the program files are copied, and then click **Next** to complete the installation.

Configure the Rendition Services Server

After you install the server software, you use Vista Configuration to configure the server to:

- access a Microsoft SQL Server 2000 or Oracle 8i database
- recognize the business services you use
- recognize the authoring applications you use

Vista Configuration Overview

The following tabs provide you access to tasks for configuring databases and authoring applications, Loading business services information and registering your CS repository.

General Tab

This tab provides general information about the machine on which the server application is running.

Database Configuration Tab

This tab lets you either create a new database from scratch or connect to an existing database. The server can access SQL or Oracle.

Domain Configuration Tab

This tab lets you load default business services contained in a pre-defined configuration file. After you configure the business services, you must install the license for your copy of the server software. The license determines the business services you're licensed to run.

Application Configuration Tab

The Application Configuration tab lets you configure the authoring application you'll use to create source documents.

FileNet Tab

The FileNet tab lets you locate the srs3mapi.dll and register your CS repository information. The repository information needs to be defined in order to complete the rendering process for a document.

Once you've completed the server configuration, you can use the Vista Administration application to verify content types. Refer to [“Configure Server Content Types” on page 680](#) for more information on verifying content types.

To open and close Vista Configuration

1. In Windows, select **Start > Program > Lquent > Vista Configuration**. The Vista Configuration application opens and the General tab displays information about the operating system and your computer.
2. When you're finished configuring the server, click **Close** to close the application.

Database Configuration Tab

The server supports Microsoft SQL Server 2000 or Oracle 8i. Procedures differ slightly for each and also depends on whether the database already exists or you are creating a new database. Refer to “Microsoft SQL Database” on page 666 for the Microsoft SQL Server 2000 configuration information. “Oracle Database” on page 667 contains the Oracle configuration information.

Notes

- Using an Access database is not supported in this release.
- You must configure the database before you can configure the Rendition Services domain.

Microsoft SQL Database

To configure a SQL database

1. In Vista Configuration, click the Database Configuration tab.
2. Under Database Configuration, select **SQL Server**.
3. In the Machine Name box, type the name of the machine on which the SQL server is located.
4. In the Catalog (database) box, type the name of the database you are creating or to which you are connecting.

CAUTION Do not enter the name of a CS library.

5. In the User Name box, type a name of the user who will access the database. The user must have Administrative rights.
6. In the Password box, type the user's password.

Note If you're connecting to an existing database, skip to step 9.

7. Click **Create New Database** to create the database. A message indicates whether the database is created.

CAUTION If the database already exists, a message will prompt you to overwrite the database. Doing so will overwrite all records stored in the database.

8. Click **OK**.
9. Click **Test Connection** to verify the database exists and that you can connect to it. A message indicates whether the connection is successful.
10. Click **OK**.

Troubleshooting the SQL Database Configuration

- If you can't create a new database, contact your administrator to see if you have access and/or a valid user name and password. Once you have the correct information re-enter it and click the Create New Database button again.
- If the test connection fails, check all the information you typed for inconsistencies or errors and try to connect again.

Oracle Database

Notes

- You must use the Oracle ODBC driver not the Microsoft Oracle driver.
- Refer to “[Prepare a Server](#)” on page 662 for more information on running the Oracle-repository.txt script before connecting to your Oracle database.
- Tablespace initial size is set to 50MB. The script turns on the tablespace auto-extend to allow unlimited growth.

To connect to an Oracle Database

1. In Vista Configuration, click the **Database Configuration** tab.
2. Under Database Configuration, select **Oracle**.
3. In the SID box, type the name of the Oracle SID that was created using the Oracle desktop client.
4. In the User Name box, type a name of the user who was created using the Oracle-repository.txt script.
5. In the Password and Confirm Password boxes, type the user's password.
6. Click **Test Connection** to verify the database exists and that you can connect to it. A message indicates whether the connection is successful.
7. Click **OK**.

Troubleshooting and Recommendations for the Oracle Database Configuration

- If the test connection fails, check all the information you typed for inconsistencies or errors and try to connect again.
- If test connection continues to fail, contact your administrator to see if you have a valid user name and password to access the database. Once you have the correct information, re-enter it and click the Test Connection button again.
- The Oracle-repository.txt contains the following default settings:

Domain Configuration Tab

After you use Vista Configuration to connect to a SQL or Oracle database, you must install and configure the business services. To do this, you can load a configuration file that was provided with the server software. The configuration file allows you to auto generate the Vista Configuration.

The configuration file allows you to auto generate the Vista Configuration.

To access the Configure dialog box

You can access the Configure dialog box from both Vista Configuration and Vista Administration. Refer to “[Check the Status of Rendition Services Jobs](#)” on page 703 for more information on Vista Administration.

1. In Vista Configuration, click the **Domain Configuration** tab.
2. Under Configuration Options, click **Advanced**. The Configure dialog box appears.

Load Configuration

To load the Rendition Services configuration

1. In Vista Configuration, click the **Domain Configuration** tab.
2. Under Configuration Options, click **Load Config**.
3. Click **Browse** next to the Configuration file box, and then browse to **c:\Program Files\Liquent\FileNet Rendition Server\Configurations\FileNet-config.xml**.
4. Click **Open**. The path and file name appear in the Configuration file box and the available configurations are displayed.
5. Highlight the option in the Available Machine Configurations box, and click **Select**. A message box confirms the configuration is created.
6. Click **OK**.

Set Up a Business Service Manager User

To set a business service manager user

1. In Vista Configuration, click the **Domain Configuration** tab.
2. Under Configuration Options, click **Set User**.
3. In the Setup Domain User dialog box, enter the name of the domain you are configuring.
4. In the User box, enter the name of the service account user.
5. In the Password box, enter the user’s password.
6. Click **OK** in the confirmation dialog box.

Load the Server License File

Your license agreement identifies the business services you're authorized to use. You must load your license after you've configured the business services.

To configure licensing

1. In Vista Configuration, click the **Domain Configuration** tab.

2. Under Configuration options, click **Advanced**.
3. Click the Licensing tab.
4. Next to the License file box, click the ellipses button.
5. Locate and highlight the license.xml file, and then click **Open**.
6. Click **Install License**. When the license is installed, it will be displayed next to License.
7. Click **OK** to close the Configure dialog and return to Vista Administration.

Set Application Configuration Tab

Use the following procedure to associate plug-ins with the correct file extensions.

To configure application plug-ins

1. In Vista Configuration, click the **Application Configuration** tab.
2. Under Set Application Data, click **Configure**.
3. In the dialog box that opens, select Microsoft Word, Corel WordPerfect, Framemaker, or Swift View, and then click **Update**. When you're finished, click **Close** to return to the Application Configuration tab.

FileNet Tab

Configuring your FileNet login entails entering the library name and system name of the CS library to which you are connecting. You can then add the Rendition queue and WatchServer entries. These two components allow the server to access publishing jobs on the CS library.

Note It is required you assign the library name to be the same name assigned for the system name when configuring a CS library for IDM Desktop. Do not use any aliases.

To configure the FileNet login

1. In Vista Configuration, click the **FileNet** tab.
2. Under Register FileNet Libraries, click **Register**. The DMS Properties dialog box appears. The srs3mapi.dll file is already located.
3. Click the + button to open the Library Properties dialog box.
4. In the Library Name box, enter the name of the CS library.
5. In the System Name box, enter the system name of the library.
Note The names assigned to the Library Name and System Name should be the same.
6. In the Property Server box, enter the name of the CS property server.
7. Click **OK**. The information you entered appears under the Library Name, System Name, and Property Server columns.

8. Click **OK** to exit the FileNet tab.

Create the FileNet Rendition Queue and WatchServer Entries

CAUTION When using an Oracle database and creating the first queue and WatchServer entry, you must use the vistaconfig.exe program directly from the Rendition Services CD. This program is also launched automatically during the Rendition Services installation. If you do not use this program to create the queue and watch server entry, you may end up with an incomplete installation and WatchServer errors.

1. In Vista Configuration, click the **FileNet** tab.
2. Under Create Queue and Watch, click **Create**. The FileNet Rendition Queue Utility dialog box appears.
3. In the Library box, select the name of the library that will contain a rendition queue. The System and Server boxes appear with your library information.
4. In the User box, enter the CS library administrator's name.
5. In the Password box, enter the CS library administrator's password.
6. Click **Create**.
7. Click **OK** in the confirmation dialog box.
8. Repeat this process until all of the libraries that will contain a rendition queue are set up.

Add Web Access

The Add Web Access option adds COM application entries to the server to allow RenderPerfect to access CS repository folders and queues. Refer to "[RenderPerfect](#)" on page 706 for more information on using RenderPerfect.

To add web access using RenderPerfect

1. In Vista Configuration, click the **FileNet** tab.
2. Under Add Web Access, click **Add** and a dialog box appears confirming you now have web access using RenderPerfect.

Complete the Server Setup

Set the Startup Properties for the Server

You must modify the BusinessServiceServer service and WatchServer service to run as the service account user.

NOTE FileNet does not recommend that either of these services be set to start automatically.

To configure the BusinessServiceServer service and WatchServer service

1. Choose **Start > Programs > Administrative Tools > Services** to access the Services window.
2. Do either of the following:

BusinessServiceServer service

- Right-click **BusinessServiceServer** on the right side of the window, and then select **Properties**. The BusinessServiceServer Properties dialog box opens.

WatchServer service

- Right-click **WatchServer** on the right side of the window, and then select **Properties**. The WatchServer Properties dialog box opens.

3. On the General tab, click the **Startup type** arrow and select **Manual**.
4. Click **Apply**.
5. Click the Log On tab, and select the **This account** option.
6. Click the **Browse** button to access the Select User dialog box and locate the name of the service account, and then click **OK**.
7. Enter the password and confirm the password for the service account user, and then click **OK**.
A confirmation appears informing you the Rendition Services service account user is granted the Log On As a Service Right.
8. Click **OK** to close the Services window.
9. Repeat these steps until both services are configured to log on as the service account.

Set the VistaDB Share Security (Oracle only)

To set the VistaDB share security

1. In Windows Explorer, go to <Windows install drive>:\Program Files\Common Files\Liquent\.
2. Right-click the **vistaDB** folder and select **Properties** to open the vistaDB Properties dialog box.
3. In the vistaDB Properties dialog box, click the **Share** tab, and then click the **Permissions** button.

4. In the Permissions VistaDB dialog box and under Name, highlight **Everyone**, and then click **Remove**. This step maintains your server security.
5. Click **Add** to open the Select Users, Computers, or Groups dialog box.
6. Click Look in arrow and select the correct domain.
7. Select the **service account user** and **IUSER_<server> users** from the list, and then click **Add**. Both users must have read/write access to both the directory security and the share permissions.
8. Click **OK** to return to the Permissions dialog box.
9. Under Permissions, click the **Allow** box next to Full Control.
10. Click **Apply**.
11. Click the **Security** tab.
12. Click the **Add**.
13. Click the **Look In** arrow and select the server domain.
14. Select the name of the Rendition Services service account user from the list, and then click **Add**.
15. Click **OK** to return to the **Security** tab.
16. Under Permissions, click **Allow** next to Full Control.
17. Click **Apply**, and then click **OK** to close the dialog box.

Reboot and Start Services

In order for the next steps to complete properly, you must reboot the server and start both the BusinessServiceServer and the WatchServer services.

To reboot and start services

- Reboot your server(s) and start both the BusinessServiceServer service and WatchServer service. Refer to [“Complete the Server Setup” on page 671](#) for more information on the BusinessServiceServer service and WatchServer service.

Set the VistaTempDir Permissions

You must set permissions for the VistaTempDir folder so that all Rendition Services account users can write temporary files to this folder. The permissions is set to Everyone by default and you should modify this setting to maintain server security.

Note The location of VistaTempDir may vary among different machines. If you're unable to locate it, create a subdirectory called VistaTempDir under the Windows temp directory.

To set up VistaTempDir permissions

1. In Windows Explorer, locate and select the VistaTempDir folder.

2. Right-click the folder and select **Sharing**.
3. Click the **Permissions** button.
4. Under Name, highlight **Everyone**, and then click **Remove**. This step lowers your risk of potential virus attacks and maintains server security.
5. Click the **Add** button.
6. Click the **Look in** arrow and select the server domain.
7. Select the name of the Rendition Services service account user from the list and click **Add**. The name of the person appears in the box at the bottom of the dialog box.
8. Click **OK** to return to the Permissions dialog box.
9. Under Permissions, click the **Allow** box next to Full Control.
10. Click **Apply**.
11. Click the **Security** tab.
12. Click the **Add**.
13. Click the **Look In** arrow and select the server domain.
14. Select the name of the Rendition Services service account user from the list, and then click **Add**.
15. Click **OK** to return to the **Security** tab.
16. Under Permissions, click **Allow** next to Full Control.
17. Click **Apply**, and then click **OK** to close the dialog box.

Modify the Lipient PDF Printer

To modify the printer settings

1. Open your Printers window.
2. Right-click **Lipient PDF Printer** and select **Properties** to access the Lipient PDF Printer Properties dialog box.
3. Click the **Advanced** tab.
4. Turn off spooling by selecting the **Print directly to the printer** option.
5. Click the **Printing Defaults** button to access the Lipient PDF Printer Printing Defaults dialog box, and then click the **Adobe PDF Settings** tab.
6. Clear the following options:
 - **Prompt for PDF File**
 - **View Result in Acrobat**
7. Click the **Edit Conversion Settings** button.
8. Click **Auto-Rotate Pages**, and then from the **Auto-Rotate Pages** list, click **Individually**.
9. Click **Save As**, and then assign a name for the new job options you just specified.
10. Click **Save**, and then click **OK** in the Job Options dialog box to return to the Lipient PDF Printer Printing Defaults dialog box.
11. Click **OK** to return to the Lipient PDF Printing Properties dialog box, and then click **OK** to exit the dialog box.

Configure the Business Service Manager

The local administrator's group must be a member of the DCOM Default Access Permission list *if there are any entries in this list*. Use the dcomcnfg utility to verify the group is in the permission list. If the group is not in the permission list, you will get an "access denied" error and your rendering job will fail.

Reboot the server and restart your services after changing this setting. Refer to "[Reboot and Start Services](#)" on page 672 for more information.

Multi-Server Install

To install Rendition Services on multiple servers

CAUTION All current rendering jobs will be aborted when adding a server to a multi-server environment. Suspend all of your watch items on each server and allow the queue to clear before attempting to add a server to your Rendition Services domain. Refer to “[Managing Watch Items](#)” on page 713 for more information.

- On each server, log in as the service account. The service account is a Windows domain user, and is also a member of the local Administrator's group. Refer to “[Prepare a Server](#)” on page 662 for more information.
- Perform the procedure for installing a single server on each server in the multi-server configuration. Each server must be installed while logged on as the service account. Each server must point to the same catalog database in the Vista Configuration tool.
- Perform the procedure in “[Complete the Multi-Server Setup](#)” on page 675 on each Rendition Services server to finish your multi-server installation.

CAUTION When pointing additional Rendition Services servers to the same catalog database, use only the **Test Connection** button in the Vista Configuration tool. Do NOT use the **Create New Database** button. Creating a new database while specifying the same catalog name as any previously-installed servers will overwrite the installed server database entries.

Complete the Multi-Server Setup

CAUTION Keep in mind the VistaTempDir is assigned to the last installed server. Refer to “[Set the VistaTempDir Permissions](#)” on page 672 for more information. A server that contains the VistaTempDir must always be available. Rendering fails if this server is not running.

To add inputs

The inputs define the services available on each server in a multi-server Rendition Services domain.

1. Access the Configure dialog box. Refer to “[Domain Configuration Tab](#)” on page 667 for the procedure on accessing the Configure dialog box.
2. Click the Domain Repositories tab, and then double-click the name of the domain you want to access. The Edit Repository dialog box appears.
3. From the Edit Repository dialog box, click the Inputs tab, and then click + to open the Add Input dialog.
4. Click the **Type** arrow and select **tcpip**.
5. Type a port number in the Channel box. The higher the port number, the less likely you'll duplicate a port that's already in use. It's recommended you select a port value of 2500 or higher.
6. Under Service, select **distribute**, **Kpublish**, **render**, **PDF File** and **retrieve**.
7. Click **OK**. The service appears on the Inputs tab.
8. Click **OK** to close the Edit Repository dialog box and return to the Configure dialog box.
9. Click **OK** to re-configure the domain.

Completed Installation and Recommendations

Your install is complete. Use the following tips for testing your installation:

General Testing Recommendations

- Start the BusinessServiceServer and WatchServer services on the Rendition Services server.
- Submit a publishing request to the Content Services server using an IDM Desktop client.
- Open the Vista Administration tool and verify the publishing request is picked up and rendered by the server.
- Check the BusinessServerManager.log and WatchServer.log in <Windows install device>:\Program Files\Common Files\Liquent\logs for errors if the server fails to render a job or pick up a publishing request.

Multi-Server Testing Recommendations

- Only start the BusinessServiceServer and WatchServer services on the last-installed Rendition Services server. This forces the newly-installed server to pick up the publishing request.
- Submit a publishing request to the Content Services server using an IDM Desktop client.
- Open the Vista Administration tool and verify the publishing request is picked up and rendered by the server.
- Once the server publishes the document, start the services on the rest of your Rendition Services servers.
- Check the BusinessServerManager.log and WatchServer.log in <Windows install device>:\Program Files\Common Files\Liquent\logs for errors if any of the servers fail to render a job or pick up a publishing request.

Disable Options in Microsoft Word and Corel WordPerfect

Before you can create PDF renditions of Microsoft Word and Corel WordPerfect, you need to turn off several options and features to ensure smooth processing by Rendition Services. Refer to the following sections for more information on the options you need to disable.

Prepare Microsoft Word Documents

Some of these items need to be configured in the documents themselves; others need to be configured in the program.

To access options and features from Microsoft Word

1. Open Microsoft Word.
2. From the Tools menu, click **Options** to access the Options dialog box.

Set the following options in the Microsoft Word application installed on the server with Rendition Services:

- Disable the Automatic Spell Checking option.
- Disable the Typing replaces selection option.
- Disable the Drag and drop text editing option.
- Disable the Use smart cut and paste option.
- Disable the View Field Codes option.
- Disable the Update Fields option.
- Disable the Update Automatic Links at Open option.
- Disable the Automatic Save Every X Minutes option.
- Disable the Confirm Conversion at Open option.
- Disable the Macro virus protection option.
- Disable the Automatically Update Document Styles option (Templates dialog box).
- Disable the Allow fast saves option.
- Disable the Allow background saves option.

Every Microsoft Word document has certain option settings embedded in it that come from settings in the Word program used when the document was saved. If the following options were not disabled, open the document in Microsoft Word, change the options, and save the document:

- Disable all autoloading macros, including macros with dialog box interfaces.
- Disable all document passwords.

For information on any of these options, see your Microsoft Word documentation.

Rendition Services cannot process a document that has links to other documents or graphic files. Before starting the processing, disable any links in the document.

If you installed Microsoft Bookshelf along with Microsoft Word, remove the Bsh32.wll file from your Microsoft Word startup folder.

Prepare Corel WordPerfect Documents

The options and features you need to turn off need to be configured in the program. Information for WordPerfect 2002 and prior to 2002 are provided below.

Corel WordPerfect 2002

Set the following options in the Corel WordPerfect 2002 application installed on the server with Rendition Services:

- In the Environment Settings dialog box, clear the **QuickTips** option.
- In the File Settings dialog box, clear the **Timed document backup** setting and the **Use enhanced file dialogs** setting.
- The registration reminder must remain offline and not appear in a pop-up window.
- The Quicklaunch applications must be disabled.

Corel WordPerfect Prior to 2002

Set the following options in the Corel WordPerfect application installed on the server with Rendition Services:

- Enable the Save Workspace To Never option.
- Disable the Spell As You Go option.
- Disable the Timed Document Backup option.

If a WordPerfect document has password protection enabled, open the document in WordPerfect, disable the password protection, and save the document.

Rendition Services cannot process a document that has links to other documents or graphic files. Before starting the processing, disable any links in the document.

For more information, see your WordPerfect documentation.

Configure the Corel WordPerfect Printer

If you are using Corel WordPerfect, you must configure its default printer after you install Rendition Services.

To configure the WordPerfect Printer

1. Start WordPerfect.
2. Select **File > Print**.
3. Set the printer to be System Default.
4. Click **Settings > Application Default > Replace**, and then click **OK**.
5. Exit WordPerfect.

Optional Steps

Use this section to perform optional procedures that appear in the “[Installation Overview](#)” on page 660.

Configure Server Content Types

Once you've configured the server using Vista Configuration, you must verify the plug-in content types in Vista Administration are correct. You access and configure content types using Vista Administration. Refer to “[Check the Status of Rendition Services Jobs](#)” on page 703 for more information on the tasks you can do with Vista Administration.

Opening and Closing Vista Administration

1. From the Start menu, select **Programs > Ligent > Vista Administration**. The StatusViewer dialog box appears.
2. To close Vista Administration, from the File menu, click **Exit**.

Application Configuration

After you've configured the business services to run on the server, you can configure the applications (authoring programs) you'll use to render documents. The native extensions associated with each application are set by default when you install the server software. When you render a document, the server checks these settings and uses the applications associated with the native extensions in the order it finds them.

If you decide to render documents using a different application plug-in, you can change the configuration. To ensure the correct application is used, you should remove the extension from the application you don't want to use. For example, if you want to render HTML using Microsoft Word rather than Microsoft Internet Explorer, you should remove the HTM file extension from Microsoft Internet Explorer. The server will find the .htm extension associated with Microsoft Internet Explorer before it finds Microsoft Word and it will use Explorer to render a document rather than Word.

Note You can add a native extension to an application that is not supported by the server; however, you may not achieve the results you expect. Some applications may render documents differently; be sure you're rendering documents using the application that best suits the rendered format. The server will render any document as long as an application can open the document.

You must also configure the server to recognize the authoring applications installed to create source documents. You can update your application configuration before you install the server software or after you've installed the authoring application.

To modify plug-in settings

1. In Vista Configuration, click the Application Configuration tab.
2. Under Modify Plugin Settings, click **Configure**.
3. Enter the additional file extension or content type you want to associate with a specific authoring application.

4. Click **Save Change**, and then click **OK**.

Access Processing Preferences for an Application

In the Content Types dialog box, you can select the plug-in for which you want to verify the content type. The content type, plug-in, and application are displayed.

To access processing preferences

1. Open Vista Administration.
2. From the View menu, click **Configure Content Types**.
3. In the Plugins list, select the application plug-in you want to verify.
4. To change the content type or description, click **Modify**. The Modify Content Type dialog box appears.
5. Click **Rendering Options** or **Bookmark Options**.
6. Make any necessary changes, and then click **OK**.

Associate Plug-ins with Content Types

For Rendition Services to process documents properly, each document file type you plan to use must be associated with a plug-in. A plug-in's main purpose is to create an accurate rendition of the document's content and layout in PDF. Some plug-ins let you customize the rendering of the content and layout or specify which entities from a document's text to extract for creating bookmarks. If a plug-in can process these kinds of features, you can specify the options when you set up a content type for it.

Each application plug-in must be associated with a content type, and a document's content type can be read from its file extension. When Rendition Services queues a document, it uses the content type (such as txt) to determine which plug-in to use for processing that document, such as the WordPad plug-in.

When you install Rendition Services, several default content types are assigned to plug-ins. The file format extensions are the standard file system extensions used by most applications. You can change the assignments and add new content types. If you have any special file formats (unique extensions), you need to establish these content type associations manually before creating PDF renditions. Refer to [“Default and Available Plug-in Options”](#) on page 696 for a list of the available content types and plug-ins.

Choosing Plug-ins

In most cases, you can process documents using the default content type assignments. However, several plug-ins can handle multiple content types, and some content types can be handled by different plug-ins. A content type can be associated with only one plug-in at a time. If you want to change the plug-in used for a content type, delete the existing content type and then add a new definition.

Handle Non-Standard File Extensions

Rendition Services uses the file name extension (such as .txt) to determine the content type of a document. If your documents do not use the default file extension for their source program, you can create a new content type. For example, if your spreadsheets are named *.feb, you can assign the Microsoft Excel plug-in to the feb content type.

If you associate a plug-in with the wrong content type, remove that content type and then re-add it with the correct plug-in.

You can also translate (remap) special characters that may appear in a file. Refer to “Remap Special Characters” on page 698 for more information.

For a list of the default content types and their plug-ins, refer to “Default and Available Plug-in Options” on page 696.

Set Plug-in Options

The Rendition Services plug-ins serve as a format translator and entity identifier for a particular file format. Some plug-ins can activate processing preferences in the supporting application. Generally, you need to set these options for a plug-in only once, when you create the content type association.

Rendition Services uses two types of processing preferences:

- **Rendering** options affect how a document's contents are laid out on the PDF rendition pages.
- **Bookmark** options define how to identify entities within the document's text for extraction. Document entities that are extracted are made into electronic bookmarks for the PDF rendition.

When a plug-in's bookmark options are set, Rendition Services identifies and gathers entities from the document's text using the option settings. The gathered entities are used to create bookmarks and a table of contents for the PDF rendition. Both the bookmarks and the table of contents entries are page-specific, electronic hyperlinks.

If an entity is extracted that has no text, its bookmark is labeled BLANK in the PDF rendition. If there is a gap in the structure of the entities extracted, a bridging bookmark is inserted into the bookmark tree to preserve the structure. The bookmark labels also appear in the generated table of contents.

Only some of the supported plug-ins provide transformation or bookmark options. You can set plug-in options for the following software and content types:

| This plug-in | Can process these file formats |
|-------------------|-----------------------------------|
| Adobe Framemaker | FM, MI, MIF |
| Adobe Illustrator | AI, BMP, EPS, GIF, JPG, PCX, TIF |
| Adobe Pagemaker | DOC, P65 |
| ASCII | ASC, DAT, SAS, TXT, RTF |
| AutoCAD | DWG, DXF |
| Corel Draw | BMP, CDR, CMX, EPS, GIF, JPG, TIF |

| This plug-in | Can process these file formats |
|-----------------------------|---|
| Corel Presentation | BMP, CGM, GIF, HPG, PCD, PIC, SHW, TGA, TIF, WMF |
| Corel Quattro Pro | DBF, DIF, QIF, SLK, TXT, WB1, WB2, WKQ, WPD, WP!, WQ1, XLS |
| Corel WordPerfect | TXT, WKB, WP, WP5, WPD, WPG |
| Isis Draw | SKC |
| Macromedia Freehand | FH9 |
| Microsoft Excel | CSV, LAB, WKS, XLA, XLC, XLS, XLT, XLW |
| Microsoft Internet Explorer | HTM, HTML |
| Microsoft PowerPoint | BMP, PPS, PPT |
| Microsoft Project | MPP |
| Microsoft Word | ASC, DOC, MCW, MSW, RTF, TXT, WPS, WRI |
| Microsoft Visio | AF2, AF3, AI, CDR, CFL, CGM, CMX, DRW, DSF, DWF, DWG, DXF, EMF, EPS, GIF, IGS, JPG, PCT, PCX, PNG, PS, VSD, VSS, VSW, WMF |
| Microstation | DGN |
| PDF | PDF |
| QuarkXpress | ASV, QXB, QXD, QXL, QXT |
| TIFFX | TIFF (single page) |

For example, text files (*.txt) can be processed by the ASCII, Corel WordPerfect, or Microsoft Word plug-in. The default plug-in assignment for txt is MS WordPad. You can delete the default assignment and use a different plug-in.

Similarly, the Acrobat Capture, Graphic Files, and Imaging plug-ins can process *.bmp files. By default, there is no plug-in assigned to the bmp content type. You can assign any one of these plug-ins to the bmp content type.

To process HTML files, use the Microsoft Internet Explorer plug-in.

Plug-in Options and Processing Preferences

ASCII Plug-in Options

The ASCII plug-in defines how to format each page and lay out the document text from a text file. The options you set determine the look of the document because text files contain no formatting information for printing.

Rendition Services uses the values you supply for paper size, paper orientation, margins, and font size along with the ESPSascii font dimensions to calculate how many text characters can fit on a line and how many lines can fit on a page. If the calculated values are greater than your selections for the maximum number of characters per line and lines per page, Rendition Services displays an error message. You must then modify the appropriate settings until you find a combination that fits on a page.

You can set these options:

- Text layout priorities:
 - Give priority to the font size you specify.
 - Give priority to the number of characters in a line and lines on a page. Rendition Services adjusts the font size to fit within the parameters.
 - Adjust the font size to fit the longest line within margins in the document. Rendition Services does not increase the font size greater than the font size you specify.
- Paper size and page orientation. The paper size should be the same dimensions as the paper dimensions in the source document.
- Font and font size for the text. Refer to [“Font Character Sets” on page 700](#) for more information.
- Page margins.
- Recognize Fortran printer codes. When the ASCII plug-in encounters the Fortran Form Feed control character (a 1 in the first column only), it starts a new page in the rendition, and prints a blank space at the beginning of that line (replacing the 1).
- Remap special characters. Refer to [“Remap Special Characters” on page 698](#). For a list of character sets, refer to [“Font Character Sets” on page 700](#).

Note Many of the ASCII transformation options interact with one another. Refer to [“How the ASCII Options Interact” on page 685](#) for details.

To access ASCII options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application” on page 681](#).
2. Choose **Rendering Options**.
3. Under Process, select a process method.
4. Use the Paper Size arrow to select the paper size for the rendition.
5. Use the Orientation arrow to select the page orientation.
6. In the appropriate boxes, type the left, right, top, and bottom margins (in inches). Use decimals for fractional sizes.
7. Use the Font arrow to select a font character set.
8. Use the Font Size arrow to select a font size.
9. Under Char Count, do any of the following:

| To | Do this |
|--|--------------------------------------|
| Set the maximum number of characters for a line. | Select Max. Characters/Line . |
| Set the maximum number of lines for a page. | Select Max. Lines/Page . |

10. To recognize Fortran printer codes when placing text on a page, select **Use Fortran Form Feed Control**.

11. Remap any special characters.

12. Choose **OK**.

How the ASCII Options Interact

When setting the transformation options for the ASCII plug-in, be aware that the options can interact with one another. Many of the options are interrelated: a change in one option may affect other options.

For example, note that the following ASCII options may be incompatible in the ESPSascii font:

- Process: Fit Char Count To Font Size
- Paper Size: US Letter
- Orientation: Portrait
- Left Margin: 0.5
- Right Margin: 0.5
- Top Margin: 0.5
- Bottom Margin: 0.5
- Font Size: 9
- Max Characters/Line: 100
- Max Lines/Page: 85

At 9 points, the ESPSascii font can fit only 93 characters per line and 80 lines per page within the specified margins on portrait, US Letter paper. One of the following adjustments may resolve this incompatibility:

- Choose another font size (such as 8 point, which can fit 105 characters per line and 90 lines per page within the margins).
- Set fewer maximum characters per line at a 9-point font size (anything equal to or less than 93 characters) and fewer maximum lines per page (anything equal to or less than 80 lines).
- Reduce the margins to values that permit the maximum characters and lines to fit on the page at that font size.

- Change the paper size or orientation to provide more space. (This option may not solve the problem entirely. Changing the paper size provides more length but reduces the available width; changing the orientation provides more width but reduces the available length.)

As a rule, Rendition Services always assumes that the paper size, orientation, and margin parameters are what you want and are not to be changed. This means that the font size and maximum characters per line and lines per page are manipulated to resolve a conflict. Remember to leave room in the margins for any headers and footers you have specified in a master page.

Many of these issues are handled automatically by Rendition Services when you use any of the process options. The following options are available:

- **Fit Char Count To Font Size** assumes that the font size is the critical factor in properly processing a document. Rendition Services displays the maximum number of characters per line and lines per page for the font size you have selected. These “limit” numbers are displayed in brackets before the appropriate text box. If the Char Count values that you entered do not accommodate the selected font size within the specified page parameters, the plug-in reduces the Char Count values to accommodate the font size. You should always check that the Char Count values you have entered are within the displayed limits.
- **Fit Font Size To Char Count** assumes the characters-per-line and lines-per-page parameters are critical factors in properly processing a document. If the font size you selected will not accommodate the Char Count values within the specified page parameters, the plug-in reduces the font size to a size that will accommodate the Char Count values. If the font size you specified does fit with room to spare, Rendition Services does not increase the font size; it uses the font size you specified.
- **Fit Font Size To Document** has the plug-in parse the document text before processing it. The plug-in finds the longest line in the document text, and sets the font size to accommodate this length within the specified page parameters. It ignores any Char Count values. Then the plug-in begins to process the document using the calculated font size.

You can use any of these process options, and you can modify any parameter at any time. If you enter an invalid combination of parameters, an error message is generated.

AutoCAD Plug-in Options

The AutoCAD plug-in processes AutoCAD engineering drawing files into PDF pages. You can select the paper size and orientation for the rendition pages. Each paper layout (not the model) in a drawing is scaled proportionally onto an individual page; the drawing is not rotated.

To access AutoCAD options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application”](#) on page 681.
2. Choose **Rendering Options**.
3. For Orientation, select the orientation for the rendition pages.
4. Use the Paper Size arrow to select the page size.
5. Choose **OK**.

CorelDRAW Plug-in Options

You can select a portrait or landscape page orientation for the rendition of a CorelDRAW image file.

To access CorelDRAW options

1. Access the processing preferences. Refer to “[Access Processing Preferences for an Application](#)” on page 681.
2. Choose **Rendering Options**.
3. Under Page Orientation, select a page orientation for the rendition.
4. Choose **OK**.

Corel Presentations Plug-in Options

The Corel Presentations plug-in processes presentation slides into PDF pages. Each slide is placed on an individual page. The slide background, title, body text and images are included in the rendition. You can have images printed in color or black and white, and you can print slides in the rendition in reverse order.

To access Corel Presentations options

1. Access the processing preferences. Refer to “[Access Processing Preferences for an Application](#)” on page 681.
2. Choose **Rendering Options**.
3. Do one of the following:
 - To print images in black-and-white, select the **Print black/white image(s)** check box.
 - To print slides beginning with the last page and ending with the first page, select the **Print in reverse order** check box.
4. Choose **OK**.

Corel Quattro Pro Plug-in Options

The Quattro Pro plug-in processes files (called notebooks) created in Quattro Pro in PDF format. You can have the rendition contain all sheets in the notebook or only the active sheet (the sheet displayed when the file is first opened), and you can select a page orientation.

To access Corel Quattro Pro options

1. Access the processing preferences. Refer to “[Access Processing Preferences for an Application](#)” on page 681.
2. Choose **Rendering Options**.
3. Under Spreadsheet, select whether to include only the active sheet or all sheets in the notebook.
4. Under Orientation, select a page orientation for the rendition.
5. Choose **OK**.

Corel WordPerfect Plug-in Options

For Corel WordPerfect documents, you can suppress the printing of headers and footers in the rendition. There are two Corel WordPerfect plug-ins: Corel WordPerfect 8 interfaces with WordPerfect version 8.0 and Corel WordPerfect interfaces with WordPerfect 9.0. You should use the plug-in that can interface with the version of WordPerfect you have installed on Generator.

You can also choose how the WordPerfect plug-in should identify document entities to extract and make into bookmarks. You can extract any document entities marked with a TOC field code. If you clear this option, no entities will be extracted. For more information on field codes, see your WordPerfect documentation.

To access Corel WordPerfect options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application”](#) on page 681.
2. Do one of the following:
 - Choose **Rendering Options** if you want to suppress the printing of headers and footers. Select the **Suppress Headers And Footers** check box.
 - Choose **Bookmark Options** if you want to extract document entities and make them into bookmarks. Select the **Extract TOC** check box.
3. Choose **OK**.

Macromedia FreeHand Plug-in Options

The Macromedia FreeHand plug-in can scale an image larger or smaller than actual size or print it at actual size in the rendition.

To access Macromedia FreeHand options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application”](#) on page 681.
2. Select **Rendering Options**.
3. Under Scale, do one of the following:

| To | Do this |
|--|---|
| Enlarge or reduce the image by a percentage. | Select the % box, and type a value between 1 and 999, where 100 is the actual size. |
| Scale the image proportionally to fit the width of the page. | Select Fit on paper . |

4. Choose **OK**.

Microsoft PowerPoint Plug-in Options

The Microsoft PowerPoint plug-in processes Microsoft PowerPoint presentation slides into PDF pages. Each slide is placed on an individual page. The slide background, title, body text and images are included in the rendition. You can define how to format each slide on a page.

The Microsoft PowerPoint plug-in can also extract the title of each slide in a PowerPoint presentation to make bookmarks for the rendition. You can specify a list of characters (using character codes) to be replaced or removed from all slide titles. To see a list of character codes, see the appropriate font character set in Appendix A.

For more information about replacing special characters, see “[Remap Special Characters](#)” on page 698.

To access Microsoft PowerPoint options

1. Access the processing preferences. Refer to “[Access Processing Preferences for an Application](#)” on page 681.
2. Do one of the following:
 - Select **Rendering Options** and do any of the following:

| To | Do this |
|--|--|
| Print a slide's text, images and notes in black and white. | Select the Black & White check box. |
| Enlarge text and images proportionally to fit the width of the printable area. | Select the Scale To Fit check box. |
| Print a border around the slide's contents. | Select the Frame Slides check box. |
| Print any associated notes below a slide's contents. | Select the Print Notes check box. |
| Print any slides as handouts, where multiple slides appear on one page. | Select the Handouts check box. |

- Select **Bookmark Options** and do any of the following:

| To | Do this |
|-----------------------------------|--|
| Add a character to the list. | <ul style="list-style-type: none"> In the Special Characters box, type the code for the character you want to replace or not print. In the Replacement Character box, type a code for the character you want to print or type a zero (0) to remove the character from the title. Choose Add. |
| Remove a character from the list. | Select a character code in the list, and choose Remove . |

3. Choose **OK**.

Microsoft Project Plug-in Options

The Microsoft Project plug-in processes project information such as tasks, resources, costs and progress created and stored in a Microsoft Project file. The project information is stored in views. The view format you select determines how the project information appears in the rendition. You can select one or more views. The plug-in does not process information in any of the Microsoft Project report formats.

You can also apply some of the Microsoft Project print options to the view.

For more information on the view formats and printing options, see your Microsoft Project documentation.

To access Microsoft PowerPoint options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application”](#) on page 681.
2. Select **Rendering Options**.
3. To turn off the printing of headers and footers, select the Hide Headers and Footers. This option applies only to Microsoft Project 95 files.
4. To scale the text to fit the information on one page width, select the **One Page Wide** check box.
5. To begin a new page at each page break inserted into the file, select the **Manual Page Breaks** check box.
6. Under Time Scale, select the range of data to include in the rendition.
7. If you selected Dates From, type the starting and ending dates.
8. Under Views, select the view format(s) for the rendition.
9. Choose **OK**.

Microsoft Word Plug-in Options

Extensive rendering options and bookmark options are available for the Microsoft Word plug-in options. This topic is divided into two sections to make it easier for you to find the rendering or bookmark information you need.

Note Before processing Microsoft Word documents, you may need to set specific options in the Word program on the server. Refer to [“Prepare Microsoft Word Documents” on page 677](#) for more information.

Rendering Options

You can suppress headers or footers, and preserve Word Web links, internal cross-references, table of contents links and cross document links in a Microsoft Word document. You can have line numbers printed on every page of a rendition and specify how the line numbering is formatted. Links are displayed as colored text. You can also specify the color of any rendered links in the original source document.

External document links must be absolute paths and the target document must not be moved or else the link will be broken. External document links are only supported for file system documents. When selected, only text is suppressed in a header or footer; graphics or drawing objects, such as lines, are not suppressed during rendering. If text is highlighted in a header or footer, the highlight is not suppressed during rendering.

Internal Microsoft Word links are preserved using the Page Layout view. This view ensures that all links are targeted to the proper page. Whenever the view is changed, however, Microsoft Word may automatically process any field codes in the document. Sometimes, the result is that the link targets are not properly defined and the links in the rendition may not jump to the proper destination. Therefore, if you want to preserve any links or cross references in a Microsoft Word document, it is recommended that you also select the Update Field Codes option. This will ensure that the field codes are updated before the rendition is created.

For more information about Word hyperlinks and line numbering, see your Microsoft Word documentation

To access Microsoft Word rendering options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application” on page 681](#).
2. Select **Rendering Options**.
3. To turn off the printing of headers, select the **Suppress Headers** check box.
4. To turn off the printing of footers, select the **Suppress Footers** check box.
5. To update all field codes in the document, select the **Update Field Codes** check box.
6. Under Word Internal Links, do any of the following:

| To | Do this |
|---|--|
| Include web hyperlinks in the rendition. | Select the Preserve Word Web Links check box |
| Include cross-reference links in the rendition. | Select the Preserve Internal XRefs check box. |

| To | Do this |
|--|--|
| Include table of contents links in the rendition. | Select the Preserve TOC Links check box. |
| Include links to external documents in the rendition. | Select the Preserve Cross Document Links check box. |
| Preserve font styles in MS Word or UltraLink cross-references. | Select the Preserve Word Font Formatting check box. |
| Specify the color of all preserved links. | Use the Link Color arrow to select the color. |

7. To add line numbers on all pages, select the Add Line Numbering check box.

8. Under Line Numbers, do the following:

| To | Do this |
|---|---|
| Set the number for the first line. | In the Start at box, type or select a number. |
| Set the distance from the text the line numbers are placed. | In the From text box, type or select a number (in inches) or select Auto . Auto uses Microsoft Word's default setting. |
| Specify how the line numbers should increment. | In the Count by box, type or select a number. |

9. Under Numbering, select how line numbers should be applied.

10. Choose **OK**.

Bookmark Options

The Microsoft Word plug-in can also extract entities from a document's text using either the styles in a document's template or an auto-generated table of contents.

If you use document templates, you can specify a list of styles that the plug-in uses to identify document entities to extract and turn into bookmarks. For each style, you specify a hierarchical level that describes how the bookmarks are related to each other, like the documents in a compound document outline. You can select a level number or one of two special level assignments:

- Subordinate is one level below the previous extract assigned a level number.
- Same is at the same level as the previous extract assigned a level number.

You can change the characters in the extracted bookmark title by using a template and selecting characters to be replaced in the title or removed from it.

For more information about replacing special characters, see [“Remap Special Characters” on page 698](#).

When extracting using document templates, you can also specify a secondary method to be used if a document's template cannot be found in the list. You can choose to either skip extracting, use another template from the list, or extract using Heading Styles or Table Entry Fields. If you use an auto-generated table of contents, figures or tables, you can extract those entities marked as either Heading styles (Heading 1, Heading 2, Heading 3 and so on) or Table Entry Fields.

Auto-generated tables of figures or tables extract any text created using Word's Insert Caption function. Tables of tables and tables of figures cannot be extracted if you are using Microsoft Word 95 on your Generator computers.

You can specify the identifiers to be used when extracting using Table Entry Fields. Each identifier is usually a single character between A and Z. In CoreDossier, you can specify more than one identifier, and the table of contents generated will contain all entities marked by Table Entry Fields with the designated identifiers. The table of contents entries are then extracted and turned into bookmarks.

You can use any combination of these extraction options. The final set of bookmarks will contain all the entities extracted from each method in the order you specified. To distinguish between multiple sets of extracts gathered by different methods, you can specify the text that should precede each set of extracts.

To access Microsoft Word bookmark options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application” on page 681](#).
2. Select **Bookmark Options**.
3. Under Extraction By, select the method(s) for identifying entities to be turned into bookmarks.
4. If you selected more than one method, specify the order of extraction. Do any of the following:

| To | Do this |
|----------------------------------|--|
| Move a method up one position. | Select a method, and then click the blue up arrow. |
| Move a method down one position. | Select a method, and then click the blue down arrow. |

5. To set Template Style options, select **Template Styles**, choose **Options**, and do any of the following:

| To | Do this |
|--------------------|---|
| Add a template. | Choose Add , select a template, and choose Open . |
| Remove a template. | Select a template, and choose Delete . |

| To | Do this |
|---|---|
| Separate multiple sets of extracts. | <ul style="list-style-type: none"> • Select the Precede Extracts With check box. • Type a label in the box. This text is used as a heading to each set of extracts produced by extraction. |
| Select a secondary method. | <p>Under If Template Not Found, select a secondary method.</p> <ul style="list-style-type: none"> • If you choose Use, use the Use arrow to select another template. • If you choose Use Table Of Contents, choose Options, specify the method options, and choose OK. |
| Under If Template Not Found, select a secondary method. | <ul style="list-style-type: none"> • If you choose Use, use the Use arrow to select another template. • If you choose Use Table Of Contents, choose Options, specify the method options, and choose OK. |
| Add a style to the list. | <ul style="list-style-type: none"> • Select a template, and choose Modify. • Under List of Styles, select a style, and click >. • Use the Level arrow to select a hierarchical level. • To remap characters, choose Special Char. • Choose OK. |
| Remove a style to the list. | <ul style="list-style-type: none"> • Select a template, and choose Modify. • Under Extract Following Styles, select a style, and click <. • Choose OK. |

6. To set Table of Contents, Figure or Table options, select the appropriate method, choose **Options**, and do any of the following:

| To | Do this |
|-------------------------------------|---|
| Extract using Heading styles. | <ul style="list-style-type: none"> • Select the Use Heading Styles check box. • In the Show levels up to box, type or select the number of levels that CoreDossier should gather. • Choose OK. |
| Extract using Table Entry Fields. | <ul style="list-style-type: none"> • Select the Use Table Entry Fields check box. • In the Type identifier for Table Entry Fields box, type one or more Table Entry Field codes to be extracted. Separate multiple codes with a colon (:). • Choose OK. |
| Separate multiple sets of extracts. | <ul style="list-style-type: none"> • Select the Precede Extracts With check box. • Type a label in the box. This text is used as a heading to each set of extracts produced by extraction. |

7. Choose **OK**.

MicroStation Plug-in Options

The MicroStation plug-in renders MicroStation PowerScope drawing files into PDF pages. You can select the paper size for the rendition pages

To access MicroStation options

1. Access the processing preferences. Refer to [“Access Processing Preferences for an Application”](#) on page 681.
2. Choose **Rendering Options**. The Microstation Transformation Options dialog box appears.
3. Under Paper Size, select the page size for the PDF rendering of a drawing.
4. Choose **OK**.

TIFF Plug-in Options

The TIFF plug-in converts scanned TIFF image files into PDF. The option settings describe the page layout and how the image is scaled in the PDF rendition. You can specify the paper size, orientation and margins of the PDF rendition's page. It is recommended that the paper settings match the intended output paper size for your publication. You can also specify how the image is scaled on the page.

To access TIFF options

1. Access the processing preferences. Refer to “Access Processing Preferences for an Application” on page 681.
2. Choose **Rendering Options**.
3. Use the Paper Size arrow to select a paper size.
4. Under Orientation, select a paper orientation.
5. Under Page Setup, type the desired page margins (in inches) in the appropriate boxes.

Note Use decimals for fractional sizes.

6. Do any of the following:

| To | Do this |
|--|---|
| Scale image proportionally to fit the width of the page. | Use the Scaling arrow to select Scale To Width . |
| Stretch the image to fill the page. | Use the Scaling arrow to select Scale To Page . |

7. Choose **OK**.

Default and Available Plug-in Options

The following tables list the available content types and their associated plug-ins. The first table lists the default content types and their associated plug-ins. The second table lists the additional content types that are available to assign to plug-ins.

The **Type** column lists the file extension for the content type. The **Plug-in** column lists the associated plug-in. The **Options** columns indicates whether you can set Transformation (T) or Bookmark (B) options for that plug-in.

Default Content Type and Plug-in Options

| Content Type | Plug-in | Options |
|--------------|-------------------|---------|
| ai | Adobe Illustrator | |

| Content Type | Plug-in | Options |
|--------------|----------------------|---------|
| cdr | Corel Draw | |
| doc | MS Word | T, B |
| dwg | AutoCAD | |
| fh7 | Macromedia Freehand | T |
| fm | Adobe FrameMaker | |
| htm | MS Internet Explorer | |
| html | MS Internet Explorer | |
| mpp | MS Project | T |
| p65 | Adobe PageMaker | |
| pdf | PDF | |
| ppt | MS PowerPoint | T, B |
| psd | Adobe PhotoShop | |
| qxd | QuarkXpress | |
| rtf | MS WordPad | |
| shw | Corel Presentations | T |
| vsd | MS Visio | |
| wb3 | Corel Quattro Pro | T |
| wpd | Corel WordPerfect | T, B |
| xls | MS Excel | |

Additional Content Type and Plug-in Options

| Content Type | Plug-in | Options |
|--------------|---------------------------------|---------|
| cdr | Corel Draw | T |
| htm, html | MS Internet Explorer MS Word | |
| rtf | ASCII MS Word | T T |
| skc | MDL ISIS/Draw Pro | |
| txt | ASCII | T |

Notes

- You can assign other content types (file extensions) to the existing plug-ins. The defaults are assigned only when you install Rendition Services.
- To use a plug-in, it must be associated with a content type. A content type can be associated with only one plug-in, but a plug-in can be used for more than one content type. If you want to change the plug-in used for a content type, delete the existing content type and then add a new definition.

Remap Special Characters

In several of the plug-ins, you can translate (remap) special characters that may appear in a file, such as Greek characters or scientific symbols. The plug-ins can handle octal, hexadecimal and decimal character codes. You build a library of commonly used special characters and assign another character to replace the special character in the rendition. You can also strip-out any character so that it will not appear in the rendition.

For example, to replace any tab found in a document extract's title with a space, type the control code for a tab (`\t`) in the Special Character box and then type the character code for a space from the appropriate font character set in the Replacement Character box. When the plug-in extracts the entities, it replaces all tabs with a space in the extract titles.

Special characters defined in a plug-in affect either the text of a document's pages in the rendition or the bookmark labels in a rendition. Only the ASCII plug-in provides the ability to select which font character set to use (ESPSascii, ESPSansi or Courier); all other plug-ins use the ESPSansi character set.

The font character sets in Appendix B list the decimal character codes for the fonts available in CoreDossier. When entering special characters, typing the character code number designates it as a decimal character code.

Note Rendition Services recognizes only three character sets. The character mapping differs between fonts. For example, the character beta (?) is code number 223 in ANSI and code number 225 in ASCII. The font set you use for remapping special characters must be the same font that was used to create the document. Refer to “[Font Character Sets](#)” on page 700 for more information.

The plug-ins can handle octal, hexadecimal, and decimal character codes. Use the following format to enter the character codes:

| Format | Type | Example |
|--------------|---------|---------|
| <i>nnn</i> | Decimal | 032 |
| <i>\nnn</i> | Octal | \040 |
| <i>0xnnn</i> | Hex | 0x020 |

You can also use the following special octal codes for some characters:

| Control | Code |
|------------------|------|
| alert | \a |
| backslash | \\ |
| backspace | \b |
| carriage return | \r |
| double quote | \" |
| form feed | \f |
| integer 0 (zero) | \0 |
| newline | \n |
| question mark | \? |
| single quote | \' |
| tab | \t |
| vertical tab | \v |

To remap special characters

1. From a plug-in, choose **Special Char**.
2. Do any of the following:

| To | Do this |
|-----------------------------------|--|
| Add a character to the list. | <ul style="list-style-type: none"> • In the Special Characters box, type the code for the character you want to replace or not print. • In the Replacement Character box, type a code for the character you want to print or type a zero (0) to strip-out the character. • Choose Add. |
| Remove a character from the list. | Select a character code in the list, and choose Remove . |

3. Choose **OK**.

Font Character Sets

This topic lists the common decimal character codes in the fonts used by some Rendition Services plug-ins. You can use these lists when you are choosing special characters to be stripped out or replaced in any of the plug-ins.

The Courier and ESPSansi fonts contain the standard ANSI character set. The ESPSascii font contains the standard ASCII character set. Use the appropriate font for processing text files. Rendition Services uses the ESPSansi font for special character replacement in all other plug-ins.

For special language or graphical characters beyond the standard ANSI (Latin 1) code page, see your system administrator for a list of the codes specific to the ANSI code pages used in your organization.

Values that appear with an empty box are not supported.

ESPSascii Font

| | | | | | | | | | | | | | |
|----|----|----|----|-----|---|-----|---|-----|---|-----|---|-----|-------|
| 0 | | 36 | \$ | 73 | I | 110 | n | 147 | ô | 184 | ¶ | 221 | █ |
| 1 | ☺ | 37 | % | 74 | J | 111 | o | 148 | ö | 185 | ¶ | 222 | █ |
| 2 | ☹ | 38 | & | 75 | K | 112 | p | 149 | ò | 186 | ¶ | 223 | █ |
| 3 | ♥ | 39 | ' | 76 | L | 113 | q | 150 | û | 187 | ¶ | 224 | α |
| 4 | ♦ | 40 | (| 77 | M | 114 | r | 151 | ù | 188 | ¶ | 225 | β |
| 5 | ♣ | 41 |) | 78 | N | 115 | s | 152 | ÿ | 189 | ¶ | 226 | Γ |
| 6 | ♠ | 42 | * | 79 | O | 116 | t | 153 | ÿ | 190 | ¶ | 227 | π |
| 7 | • | 43 | + | 80 | P | 117 | u | 154 | Û | 191 | ¶ | 228 | Σ |
| 8 | ▣ | 44 | , | 81 | Q | 118 | v | 155 | ç | 192 | ¶ | 229 | σ |
| 9 | ◦ | 45 | - | 82 | R | 119 | w | 156 | £ | 193 | ¶ | 230 | μ |
| 10 | ▣ | 46 | . | 83 | S | 120 | x | 157 | ¥ | 194 | ¶ | 231 | τ |
| 11 | ø | 47 | / | 84 | T | 121 | y | 158 | € | 195 | ¶ | 232 | φ |
| 12 | ♀ | 48 | 0 | 85 | U | 122 | z | 159 | f | 196 | ¶ | 233 | ⊖ |
| 13 | ♪ | 49 | 1 | 86 | V | 123 | { | 160 | à | 197 | ¶ | 234 | Ω |
| 14 | ♪ | 50 | 2 | 87 | W | 124 | | 161 | í | 198 | ¶ | 235 | ø |
| 15 | ★ | 51 | 3 | 88 | X | 125 | } | 162 | ó | 199 | ¶ | 236 | ∞ |
| 16 | ▶ | 52 | 4 | 89 | Y | 126 | ~ | 163 | ú | 200 | ¶ | 237 | ∅ |
| 17 | ◀ | 53 | 5 | 90 | Z | 127 | • | 164 | ñ | 201 | ¶ | 238 | ε |
| 18 | ↕ | 54 | 6 | 91 | [| 128 | Ç | 165 | Ñ | 202 | ¶ | 239 | ∩ |
| 19 | !! | 55 | 7 | 92 | \ | 129 | ü | 166 | ª | 203 | ¶ | 240 | ≡ |
| 20 | ¶ | 56 | 8 | 93 |] | 130 | é | 167 | º | 204 | ¶ | 241 | ± |
| 21 | § | 57 | 9 | 94 | ^ | 131 | â | 168 | ¿ | 205 | ¶ | 242 | ≥ |
| 22 | ■ | 58 | : | 95 | ~ | 132 | ä | 169 | ƒ | 206 | ¶ | 243 | ≤ |
| 23 | ↕ | 59 | ; | 96 | ` | 133 | à | 170 | ƒ | 207 | ¶ | 244 | [|
| 24 | ↑ | 60 | < | 97 | a | 134 | á | 171 | ½ | 208 | ¶ | 245 |] |
| 25 | ↓ | 61 | = | 98 | b | 135 | ç | 172 | ¼ | 209 | ¶ | 246 | ÷ |
| 26 | → | 62 | > | 99 | c | 136 | ê | 173 | i | 210 | ¶ | 247 | ≈ |
| 27 | ← | 63 | ? | 100 | d | 137 | ë | 174 | « | 211 | ¶ | 248 | * |
| 28 | └ | 64 | @ | 101 | e | 138 | è | 175 | » | 212 | ¶ | 249 | • |
| 29 | ↔ | 65 | A | 102 | f | 139 | ÿ | 176 | ☼ | 213 | ¶ | 250 | · |
| 30 | ▲ | 66 | B | 103 | g | 140 | î | 177 | ☼ | 214 | ¶ | 251 | √ |
| 31 | ▼ | 67 | C | 104 | h | 141 | ï | 178 | ☼ | 215 | ¶ | 252 | n |
| 32 | | 68 | D | 105 | i | 142 | Ï | 179 | | 216 | ¶ | 253 | z |
| 33 | ! | 69 | E | 106 | j | 143 | Ä | 180 | † | 217 | ¶ | 254 | █ |
| 34 | " | 70 | F | 107 | k | 144 | É | 181 | ‡ | 218 | ¶ | 255 | □ |
| 35 | # | 71 | G | 108 | l | 145 | æ | 182 | ‡ | 219 | ¶ | | |
| | | 72 | H | 109 | m | 146 | Æ | 183 | ¶ | 220 | █ | | |
| | | | | | | | | | | | | | ASCII |

ESPSansi Font

| | | | | | | | | | | | | | |
|----|----|----|---|-----|---|-----|-----|-----|----|-----|----|-----|---|
| 0 | □ | 37 | % | 74 | J | 111 | o | 148 | ” | 185 | ı | 222 | Ɔ |
| 1 | □ | 38 | & | 75 | K | 112 | p | 149 | • | 186 | ° | 223 | ß |
| 2 | □ | 39 | ' | 76 | L | 113 | q | 150 | – | 187 | » | 224 | à |
| 3 | □ | 40 | (| 77 | M | 114 | r | 151 | — | 188 | ¼ | 225 | á |
| 4 | □ | 41 |) | 78 | N | 115 | s | 152 | ~ | 189 | ½ | 226 | â |
| 5 | □ | 42 | * | 79 | O | 116 | t | 153 | ™ | 190 | ¾ | 227 | ã |
| 6 | □ | 43 | + | 80 | P | 117 | u | 154 | š | 191 | ı̇ | 228 | ä |
| 7 | □ | 44 | , | 81 | Q | 118 | v | 155 | › | 192 | À | 229 | å |
| 8 | □ | 45 | - | 82 | R | 119 | w | 156 | œ | 193 | Á | 230 | æ |
| 9 | □ | 46 | . | 83 | S | 120 | x | 157 | □ | 194 | Â | 231 | ç |
| 10 | □ | 47 | / | 84 | T | 121 | y | 158 | □ | 195 | Ã | 232 | è |
| 11 | □ | 48 | 0 | 85 | U | 122 | z | 159 | ÿ | 196 | Ä | 233 | é |
| 12 | □ | 49 | 1 | 86 | V | 123 | { | 160 | | 197 | Å | 234 | ê |
| 13 | □ | 50 | 2 | 87 | W | 124 | | 161 | ı̇ | 198 | Æ | 235 | ë |
| 14 | □ | 51 | 3 | 88 | X | 125 | } | 162 | ϕ | 199 | Ç | 236 | ì |
| 15 | □ | 52 | 4 | 89 | Y | 126 | ~ | 163 | £ | 200 | È | 237 | í |
| 16 | □ | 53 | 5 | 90 | Z | 127 | □ | 164 | ¤ | 201 | É | 238 | î |
| 17 | □ | 54 | 6 | 91 | [| 128 | □ | 165 | ¥ | 202 | Ê | 239 | ï |
| 18 | □ | 55 | 7 | 92 | \ | 129 | □ | 166 | ı̇ | 203 | Ë | 240 | ð |
| 19 | □ | 56 | 8 | 93 |] | 130 | , | 167 | § | 204 | Ì | 241 | ñ |
| 20 | □ | 57 | 9 | 94 | ^ | 131 | f | 168 | ” | 205 | Í | 242 | ò |
| 21 | □ | 58 | : | 95 | _ | 132 | ” | 169 | © | 206 | Î | 243 | ó |
| 22 | □ | 59 | ; | 96 | ` | 133 | ... | 170 | ª | 207 | Ï | 244 | ô |
| 23 | □ | 60 | < | 97 | a | 134 | † | 171 | « | 208 | Ð | 245 | õ |
| 24 | □ | 61 | = | 98 | b | 135 | ‡ | 172 | ¬ | 209 | Ñ | 246 | ö |
| 25 | □ | 62 | > | 99 | c | 136 | ^ | 173 | - | 210 | Ò | 247 | ÷ |
| 26 | □ | 63 | ? | 100 | d | 137 | ‰ | 174 | ® | 211 | Ó | 248 | ø |
| 27 | □ | 64 | @ | 101 | e | 138 | Š | 175 | - | 212 | Ô | 249 | ù |
| 28 | □ | 65 | A | 102 | f | 139 | < | 176 | ° | 213 | Õ | 250 | ú |
| 29 | □ | 66 | B | 103 | g | 140 | Œ | 177 | ± | 214 | Ö | 251 | û |
| 30 | □ | 67 | C | 104 | h | 141 | □ | 178 | ² | 215 | × | 252 | ü |
| 31 | □ | 68 | D | 105 | i | 142 | □ | 179 | ³ | 216 | Ø | 253 | ý |
| 32 | | 69 | E | 106 | j | 143 | □ | 180 | ´ | 217 | Ù | 254 | þ |
| 33 | ! | 70 | F | 107 | k | 144 | □ | 181 | μ | 218 | Ú | 255 | ÿ |
| 34 | ” | 71 | G | 108 | l | 145 | ´ | 182 | ¶ | 219 | Û | | |
| 35 | # | 72 | H | 109 | m | 146 | ´ | 183 | · | 220 | Ü | | |
| 36 | \$ | 73 | I | 110 | n | 147 | “ | 184 | , | 221 | Ý | | |

ANSI

Courier Font

| | | | | | | | | | | | | | |
|----|----|----|---|-----|---|-----|-----|-----|---|-----|---|-----|---|
| 0 | □ | 37 | % | 74 | J | 111 | o | 148 | ” | 185 | ¹ | 222 | Ɔ |
| 1 | □ | 38 | & | 75 | K | 112 | p | 149 | • | 186 | ° | 223 | β |
| 2 | □ | 39 | ' | 76 | L | 113 | q | 150 | – | 187 | » | 224 | à |
| 3 | □ | 40 | (| 77 | M | 114 | r | 151 | — | 188 | ¼ | 225 | á |
| 4 | □ | 41 |) | 78 | N | 115 | s | 152 | ~ | 189 | ½ | 226 | â |
| 5 | □ | 42 | * | 79 | O | 116 | t | 153 | ™ | 190 | ¾ | 227 | ã |
| 6 | □ | 43 | + | 80 | P | 117 | u | 154 | š | 191 | ç | 228 | ä |
| 7 | □ | 44 | , | 81 | Q | 118 | v | 155 | > | 192 | À | 229 | å |
| 8 | □ | 45 | - | 82 | R | 119 | w | 156 | œ | 193 | Á | 230 | æ |
| 9 | □ | 46 | . | 83 | S | 120 | x | 157 | □ | 194 | Â | 231 | ç |
| 10 | □ | 47 | / | 84 | T | 121 | y | 158 | □ | 195 | Ã | 232 | è |
| 11 | □ | 48 | 0 | 85 | U | 122 | z | 159 | ÿ | 196 | Ä | 233 | é |
| 12 | □ | 49 | 1 | 86 | V | 123 | { | 160 | | 197 | Å | 234 | ê |
| 13 | □ | 50 | 2 | 87 | W | 124 | | 161 | ı | 198 | Æ | 235 | ë |
| 14 | □ | 51 | 3 | 88 | X | 125 | } | 162 | ø | 199 | Ç | 236 | ì |
| 15 | □ | 52 | 4 | 89 | Y | 126 | ~ | 163 | £ | 200 | È | 237 | í |
| 16 | □ | 53 | 5 | 90 | Z | 127 | □ | 164 | ¤ | 201 | É | 238 | î |
| 17 | □ | 54 | 6 | 91 | [| 128 | □ | 165 | ¥ | 202 | Ê | 239 | ï |
| 18 | □ | 55 | 7 | 92 | \ | 129 | □ | 166 | ı | 203 | Ë | 240 | ð |
| 19 | □ | 56 | 8 | 93 |] | 130 | , | 167 | § | 204 | Ì | 241 | ñ |
| 20 | □ | 57 | 9 | 94 | ^ | 131 | f | 168 | ” | 205 | Í | 242 | ò |
| 21 | □ | 58 | : | 95 | _ | 132 | ” | 169 | © | 206 | Î | 243 | ó |
| 22 | □ | 59 | ; | 96 | ` | 133 | ... | 170 | ª | 207 | Ï | 244 | ô |
| 23 | □ | 60 | < | 97 | a | 134 | † | 171 | « | 208 | Ð | 245 | õ |
| 24 | □ | 61 | = | 98 | b | 135 | ‡ | 172 | ¬ | 209 | Ñ | 246 | ö |
| 25 | □ | 62 | > | 99 | c | 136 | ^ | 173 | - | 210 | Ò | 247 | ÷ |
| 26 | □ | 63 | ? | 100 | d | 137 | ‰ | 174 | ® | 211 | Ó | 248 | ø |
| 27 | □ | 64 | @ | 101 | e | 138 | Š | 175 | - | 212 | Ô | 249 | ù |
| 28 | □ | 65 | A | 102 | f | 139 | < | 176 | ° | 213 | Õ | 250 | ú |
| 29 | □ | 66 | B | 103 | g | 140 | Œ | 177 | ± | 214 | Ö | 251 | û |
| 30 | □ | 67 | C | 104 | h | 141 | □ | 178 | ² | 215 | × | 252 | ü |
| 31 | □ | 68 | D | 105 | i | 142 | □ | 179 | ³ | 216 | Ø | 253 | ý |
| 32 | | 69 | E | 106 | j | 143 | □ | 180 | ´ | 217 | Ù | 254 | þ |
| 33 | ! | 70 | F | 107 | k | 144 | □ | 181 | µ | 218 | Ú | 255 | ÿ |
| 34 | " | 71 | G | 108 | l | 145 | ‘ | 182 | ¶ | 219 | Û | | |
| 35 | # | 72 | H | 109 | m | 146 | ’ | 183 | · | 220 | Ü | | |
| 36 | \$ | 73 | I | 110 | n | 147 | “ | 184 | , | 221 | Ý | | |

Courier

Check the Status of Rendition Services Jobs

Use Vista Administration to view the status of the rendition jobs you are processing. You access the StatusViewer dialog box, where you can check the job queue to determine which rendition jobs are waiting to be processed and which ones are finished. The dialog box also displays the status (Completed, Failed, In Progress, and Succeeded) of each job in the queue. If the job fails, an error message also appears. You can delete one or all jobs with a status of Completed, Succeeded or Failed from the queue, but you cannot delete a job that has a status of In Progress.

Additional tasks can also be accessed from Vista Administration. Those tasks are: configure a domain, configure content types and configure CS library information.

The CS Admin Tools provide the Rendition Queue Administrator. The Rendition Queue Administrator is a tool that displays only the CS Rendition queue. You must use this tool if you are having problems rendering documents because it is used to delete documents from the CS Rendition queue. Refer to the CS Admin Tools help for more information on the Rendition Queue Administrator.

RenderPerfect is another tool you can use to monitor the system queues and schedule the processing hours. Refer to [“Control Rendition Services Jobs” on page 706](#) for more information on RenderPerfect.

To access Vista Administration

1. From the Start menu, select **Programs > Lquent > Vista Administration**. The StatusViewer dialog box appears.
2. To close Vista Administration, from the File menu, click **Exit**.

Domain Configuration from Vista Administration

Your domain configuration can be accessed from Vista Administration. You loaded your configuration file when you first installed Rendition Services. This is another way to access your configuration information.

To access your domain configuration

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the View menu, click **Configure Domain**.
3. Refer to [“Domain Configuration Tab” on page 667](#) for more information on procedures and tasks for working with your domain configuration.

Configure Content Types from Vista Administration

To access the content types

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the View menu, click **Configure Content Types**.
3. Refer to [“Access Processing Preferences for an Application” on page 681](#) for more information on using and modifying content types.

Work with Rendition Services Jobs

Once your jobs are being processed, you monitor the jobs in the StatusViewer dialog box. You can filter and refresh the jobs as well as delete and archive the jobs.

At the bottom of the StatusViewer dialog box, you can also view in detail the status and results of your Rendition Services jobs using the following four tabs: Job Status Detail, Current Results, Original Request, and Job Status Description. When you click on each tab, you can view the specific information on your Rendition Services job.

Filter Jobs

Use the Filter Jobs option to organize the jobs displayed in the StatusViewer dialog box.

To filter jobs

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the View menu, click **Filter Jobs**. The Job Filter Criteria dialog box appears.
3. You can optionally enter the Job ID, name of the Submitter, Job Name, or Product Name.
4. Select the status of the job you want to filter. The status options are: Completed, Failed, In Progress, and Succeeded.
5. Select the activity of the job you want to filter. The activity option is queue.
6. Under Date/Time, specify the date and time you want to filter the job. Do one of the following:
 - Select **None** if you do not want to filter a job using a date and time range.
 - Select **Today** if you want to filter jobs processed only today.
 - Select **Last** if you want to filter jobs processed only for a specified number of hours. Enter the number of hours.
 - Select **Date Range** if you want to filter jobs processed for a period of time. Select **From Date** or **From Time** to specify the starting date and time for filtering the jobs. Select **To Date** or **To Time** to specify the end date and time for filtering the jobs. Use the date and time arrows to modify the date and time ranges.
7. Click **OK** when you are finished.

Refresh Jobs

You can update the status of a Rendition Services job from Vista Administration by using the refresh job option. The jobs can also be refreshed automatically.

CAUTION Be aware there may be a large number of jobs to display and performance could be impacted when refreshing your jobs automatically or frequently.

To refresh jobs

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the View menu, click **Refresh Jobs**.

To auto refresh jobs

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the View menu, click **Auto Refresh**. A check mark appears next to the option and the status of your jobs will be refreshed automatically.

Delete Jobs

You can delete Rendition Services jobs from the jobs database by using the Delete Jobs option.

To delete jobs

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the Jobs menu, click **Delete Jobs**.

Archive Jobs

You can archive Rendition Services jobs from the jobs database by using the Archive Jobs option. For more information on archiving jobs, refer to [“Specify Archived Jobs” on page 714](#).

To archive jobs

1. Access Vista Administration to open the StatusViewer dialog box.
2. From the Jobs menu, click **Archive Jobs**.

Control Rendition Services Jobs

RenderPerfect

RenderPerfect is a utility that allows you to change queueing times and days. It also helps you configure Rendition Services to render files on local drives.

When you install Rendition Services, the installer sets up default WatchServer entries. These default entries allow you to render documents from the CS library publishing queue, and they are set to check that queue every 30 seconds.

If you require different or additional queue or watch folder functionality, the information provided here contains instructions for setting up a customized RenderPerfect configuration.

Using RenderPerfect

You can combine items from one or more locations (file system, CS) to watch. When you select the local file system to watch, folders are displayed on the Select Watch Item pane.

If you do not have a RenderPerfect login, you must create one using the Vista Configuration application. Refer to [“Change the RenderPerfect Login”](#) on page 716 for more information.

RenderPerfect uses the following items to monitor your Rendition Services jobs: Action Groups, Folder Groups and Watch Items. Action Groups allow you to specify actions performed when rendering documents. Folder Groups define the folders and queues the server checks. Watch items combine information from the Action Groups and Folder Groups with scheduling information to render documents to PDF.

Refer to the following sections for more information on using RenderPerfect.

To start RenderPerfect

1. From the Start menu, highlight **Programs**, point to **Liquent** and select **RenderPerfect**. The Administrator Login pane appears in your default browser.
2. In the Username and Password boxes, type the default username and password.

Note If you changed the default Username and/or Password when you configured the software, use those values to log in. For details about changing the default login, see [“Change the RenderPerfect Login”](#) on page 716 for more information.

3. Click **Log-in**. The Administrator Welcome pane appears. The activities you can perform appear as links on the QueueBuilder Menu pane on the left side of the browser window.

To exit RenderPerfect

- On the QueueBuilder Menu pane, click **Logoff RenderPerfect**. You are logged off RenderPerfect; the Administrator Login pane appears so you can log back in.

Action Groups

Action groups enable you to group together the actions to be performed when rendering documents. You can select options for rendering to PDF. Once you've selected the options and saved the Action Groups, you can manually enter additional options that aren't available as selections in the application. In addition, you can create your own action group by creating an XML ticket manually.

An XML ticket is a job ticket that defines the documents to be processed and the processes to be performed using RenderPerfect.

This section is divided into the following:

- Creating an Action Group describes how to select options and build the XML ticket automatically.
- Creating an Action Group Manually describes how to add advanced actions to an action group, and how to create an action group by entering the XML ticket manually.
- Managing Action Groups describes how to edit and delete action groups.

Creating an Action Group

To create an action group

1. On the QueueBuilder Menu pane, click **Manage Action Groups**. The Select Action Group pane appears.
2. Click **Create New Action Group**. The Create Action Group pane appears.
3. In the Action Group Name box, type a name for the new action group. The options you select under Render determine the additional options that will be available for rendering documents. See the following sections for detailed information about each rendition type and its options:
 - Creating a Render-to-PDF Action Group
 - Creating a Render-to-XML Action Group
 - Checking and Saving the Action Group

Note The Render-to-XML Action Group option is available, but it is not supported in the current release.

Creating a Render-to-PDF Action Group

To create a render-to-PDF action group

1. Under Render, select **PDF**.
2. To save a copy of the rendition in the location of the source document, select **Save as Rendition** to the right of PDF.
3. Under Publish, select **PDF Watermark** to create a watermark on the document pages.
4. In the box provided, type text for the watermark
5. In the Position box, type one of the following values to indicate where on the page to position the watermark:

| Value | Location |
|-------|---------------|
| 0 | Left edge |
| 1 | Right edge |
| 2 | Top edge |
| 3 | Bottom edge |
| 4 | Top to bottom |
| 5 | Bottom to top |
| 6 | Center |

6. In the Intensity box, type one of the following to indicate how dark the watermark should be:

| Value | Location |
|-------|-----------|
| 0 | Hollow |
| 10 | 10% black |
| 20 | 20% black |
| ... | ... |
| 100 | Black |

7. Under Security, select **PDF Security** to specify security options for the rendition.

8. Select any of the following security options:

- **Allow Printing** lets users print the PDF rendition.
- **Allow Changing** lets users edit the rendition.
- **Allow Content Copying** lets users copy content from the rendition to paste in other documents.
- **Allow Changing Comments** lets users edit Acrobat annotations in the rendition.

9. In the User Password box under the Allow options, type the password users must enter to access the rendition.

10. In the Owner Password box under the Allow options, type the password the rendition owner must enter to access the rendition.

11. As you select options for the action group, an XML ticket is created.

12. To view the XML ticket, click **Display XML**. The job ticket appears at the bottom of the Create Action Group pane.
13. When you're finished, click **OK**.

Creating a Render-to-XML Action Group

Note This feature is not currently supported in this release.

Editing the XML Ticket

As you select the options for the action group, an XML ticket is created. You can view the XML ticket and include additional actions using the Advance button on the Create Action Group pane.

To view and edit the XML ticket

1. On the QueueBuilder Menu, click **Manage Action Groups**.
2. On the Select Action Group pane, click **View/Edit** next to the action group you want to edit.
3. On the Create Action Group pane, click **Display XML**. The ticket is displayed in a text box below the options.
4. To include additional actions, select all the text in the ticket.
5. From the browser menus select **Edit > Copy**.
6. On the Create Action Group pane, click **Advanced**. The Enter XML Ticket pane appears.
7. To paste the copied ticket text, click in the text box, then select **Edit > Paste**.
8. Manually add the actions you want to include.
9. When you're satisfied with the action group, type a name for it in the Action Group Name box.
10. To prevent the action group from being edited, select **Read Only** next to the action group name.
11. Click **OK** to save the action group. A message indicates the new action group is saved.

Creating an Action Group Manually

If you prefer, you can write your own action group ticket using the Enter XML Ticket pane.

To create an action group manually

1. On the QueueBuilder Menu pane, click **Manage Action Groups**.
2. On the Select Action Group pane, click **Create New Action Group**.
3. Click **Advanced** at the bottom of the pane. The Enter XML Ticket pane appears.
4. In the Action Group Name box, type in a name for the action group (ticket) you are creating.

5. In the box provided, type the actions for the action group.

Note You can create a ticket using the Create Action Group options, click the Display XML button to display the ticket, and then copy and paste it into this box to add actions.

6. To prevent anyone from changing the action group, select **Read Only** next to the action group name at the top of the Enter XML Ticket pane.
7. When you're finished, click **OK** to save the new action group. A message indicates the new action group is saved.

Managing Action Groups

Managing action groups includes the following:

- Viewing and editing an action group
- Deleting action groups

To manage action groups

1. On the QueueBuilder Menu pane, click **Manage Action Groups**. The Select Action Group pane appears. All action groups are displayed.
2. To delete action groups, do either of the following:
 - To delete one action group, select its check box and click **Delete**.
 - To delete all action groups, select **Action Groups** and click **Delete**.
3. To view the contents of an action group or change it, click View/Edit to the right of the action group.

Note If the Read Only check box was selected when an action group was created, there is no View/Edit link for the action group; it cannot be edited.
4. The action group appears in the Edit Action Group pane where you can make changes, then click **OK** to save them.

Folder Groups

Folder Groups let you point the Rendition Services server to specified file system folders and queues. After you create all of the folder groups, you can create watch items. For details, refer to [“Watch Items” on page 712](#) for more information.

This section covers the following:

- Creating a Folder Group describes how to create a folder group.
- Selecting Folders describes how to add file system, such as CS items to a folder group.
- Managing Folder Groups describes how to view, change, and remove folder groups.

Creating a Folder Group

To create a folder group

1. On the QueueBuilder Menu pane, select **Manage Folder Groups**. The Select Folder Group pane appears.
2. Click **Create New Folder Group**. The Enter New Folder Group Name pane appears.
3. In the Folder Group Name box, type in a name for the new folder group.
4. In the Folder Group Description box, type in a short description for the new folder group.
5. To create the new folder group, click **OK**. The Select Watch Item pane appears with the name of the new folder group at the top of the pane.

Selecting Folders

You can combine folders from one or more locations to make up a folder group. When you add a folder, it's saved automatically as part of the group.

Note Before accessing CS libraries from RenderPerfect, make sure you first use Vista Configuration to configure RenderPerfect to access CS repository folders and queues. Refer to [“Add Web Access” on page 670](#) for more information.

Selecting Folders from a File System

To select folders from a file system

1. On the Select Watch Item pane, select **File System**. Its folders appear on the Select Watch Item pane.
2. Locate each folder to include in the folder group and click add to Folder Group. Each folder you select moves to the Folder Group section of the pane.
3. To add items from a CS library, click Up one level to return to the pane with the File System and CS links. Items you add to the folder group are saved automatically.

Selecting Items from CS

To select folders from CS

1. Select **CS**. The databases appear on the Select Watch Item pane.
2. Locate and click the database you want to select items from. The User Log in box appears on the Log in to DMS pane along with the name of the database you selected.
3. Type the CS username and password in the boxes provided, then click Log-in. The folders in the database appear under Select Watch Item.
4. Locate each folder to include in the folder group and click add to Folder Group. Each folder you select moves to the Folder Group section of the pane.

- To add items from another CS library, or file system, click Up one level to return to the pane with the File System and DMS links. Items you add to the folder group are saved automatically.

Assigning a Filter to a New Folder Group

Note This feature is not currently supported in this release.

Managing Folder Groups

Managing folder groups includes the following:

- Viewing and editing a folder group
- Removing a filter group assignment
- Removing a folder group

To manage a folder group

- On the QueueBuilder Menu pane, click Manage Folder Groups. The Select Folder Group pane appears. All folder groups are displayed.
- To delete folder groups, do one of the following:
 - To delete one folder group, select its check box and click **Delete**.
 - To delete all folder groups, select the **Action Groups** check box and click **Delete**.
- To view the contents of a folder group or change it, click **View/Edit**. The Folder Group pane appears.
- Do any of the following:
 - To add a filter to the group, click **Add Filter**.
 - To remove a folder from the group, click **Remove From Group**.
 - To add more folders, locate a folder then click **add to Folder Group**.
 - To remove a filter from a group, click **View Filter**, then click **Remove on the Assign Filter to Folder** pane.

When you're finished, a message indicates the settings were saved.

Watch Items

Watch items are actions you associate with specific folder and action groups. When you create a watch item, you can specify when it will run.

The watched folder functionality renders all documents inside a folder every time the watch is requested. Therefore, do not run a watch on that folder frequently. Instead run it once a day or at most every several hours.

This section is divided into the following:

- Creating a Watch Item
- Managing Watch Items

Creating a Watch Item

To create a watch item

1. On the QueueBuilder Menu pane, select **Create Watch Item**. The Edit Watch pane appears.
2. Click the Folder Groups arrow and select a folder group to associate with the watch item.
3. Click the Action Groups arrow and select an action group to associate with the watch item.
4. Under Job Options, type a name for the watch item in the Name box.
5. Under Job should run, select how often you want the job to be run.
 - Once runs the job one time only.
 - Daily runs the job each day using the options at the bottom of the pane.
 - Weekdays runs the job on the weekdays you select using the checkboxes provided.
 - Monthly runs the job once each month.
6. If you select Daily, do the following:
 - Select **Every**, type an interval, then click the select Seconds, Minutes, or Hours from the dropdown box.
 - Select **Start At**, type a time, then select **AM** or **PM**.
 - Select **Run To**, type a time, then select **AM** or **PM**.
7. Under Job should start, enter a date range during which the job should run.

Note To display a calendar from which you can select a date, click Select Date, use the arrows to select the month and year, and click a date to select it.
8. Click the close box to close the calendar window.
9. Click **OK** to save the watch item.

Managing Watch Items

Managing watch items includes the following:

- Suspending and restarting watch items
- Deleting watch items
- Editing watch items

To manage watch items

1. On the QueueBuilder Menu pane, select **View Watch List**. The Select Watch pane appears.
2. To delete the watch items, do the following:
 - To delete one watch item, select it and then click **Delete**.
 - To delete all watch items, select **Watch Items**, then click **Delete**.
3. To suspend watch items, do one of the following:
 - To suspend one watch item, select it, then click **Suspend**.
 - To suspend all watch items, select **Watches**, then click **Suspend**.
4. To edit a watch item, click **View/Edit**.

Specify Archived Jobs

When you specify a job archive machine you are allocating responsibility for the job archiving service. All jobs for all machines in a domain are stored in a Job table in the database. Whichever machine has been specified as the archive machine will periodically move Jobs from the Jobs table into an Archived Jobs table, and purge jobs from that Archive table according to the specified interval.

The default Purge Interval setting is set to 7 days. The default Archive Interval setting is set to 1 day.

Note You can view archived jobs in Vista Administration. Refer to [“Archive Jobs” on page 705](#) for more information on using this the Archive Jobs option.

Watched Folders and IDM Desktop 4.0 (or higher) Preferences

When you are using watched folders with Rendition Services 5.5, Desktop Publishing preferences are also applied on the Renditions Services server.

Notes

- The preferences are only used for the first version of a published document. Once a document has a published version, the Desktop preferences are ignored and the document will be published the same way a previous version was published.
- Whenever the Desktop preferences are modified on the Rendition Services server, you must reboot the server for the preferences to take effect.
- If you have multiple servers accessing the same queue, make sure your publishing preference settings are the same on each Rendition Services server. Each watched folder will always be published according to the same preferences.

To access the IDM Desktop Publishing Preferences

1. From the Windows Start menu, select **Programs, FileNet IDM**, and then **Configure**.
2. Select **Set default preferences for new users** or **Set preferences for currently logged on user**.
3. Click **Additional Preferences**. The IDM Preferences dialog box displays the list of categories: Documents, Directories and Files, Searching, Publishing, Viewing, Libraries, and Application Integration, and Print/Fax.
4. Select the Publishing preference.
5. Make your changes, click **OK**, and then click **Close**.

Refer to the following table that documents the Desktop publishing preferences and special notes explaining how some of the preferences affected watched folders and published documents.

| Publishing Preference | Notes |
|---|--|
| Document Options | The settings for this preference are always ignored and the watched folder published documents will be versioned and not replaced. |
| Allow changes to “Document Options” | N/A |
| Security for new published documents | This preference applies to the first published version published using a watched folder. |
| Delete published document if source is deleted | This preference applies to the first published version published using a watched folder. |
| Show “Delete published document if source is deleted” | N/A |
| Inherit source document properties | <ul style="list-style-type: none"> • This preference applies to the first published version published using a watched folder. • Always select this preference because publishing a document with a document class that contains required properties will fail if this setting is not selected. |
| Folder ID | This preference applies to the first published version published using a watched folder. |

Log Files

Rendition Services generates logs while processing jobs. The two log files that are generated are BusinessServiceManager.log and WatchManager.log.

The BusinessServiceManager.log file provides information on your Rendition Services server, such as when the server starts and stops while processing jobs. The WatchManager.log files provides information on specified watched folders. In addition, the BusinessServiceManager.log and WatchManager.log files are located in your <install device>:\Program Files\Common Files\liquent\log folder.

These log file messages are generated while processing jobs. If one of your jobs fails, you can open the log file in a text editor to help identify the problem.

To access the BusinessServiceManager.log and WatchManager.log files

1. Open Windows Explorer.
2. Locate the log folder in <install device>:\Program Files\Common Files\liquent\.
3. Double-click the log folder to access the BusinessServiceManager.log and WatchManager.log files.
4. Double-click either file to view in a text editor.

Change the RenderPerfect Login

If you plan on using RenderPerfect, you must also set up login information for it. RenderPerfect is a web interface and using it involves allowing access to CS repository folders and queues via the web. Refer to “[Add Web Access](#)” on page 670 for more information on setting RenderPerfect. For more information on using RenderPerfect, refer to “[Control Rendition Services Jobs](#)” on page 706.

To configure a RenderPerfect login

1. Access the Configure dialog box.
2. Click + to open the Add Login dialog box.
3. In the What box, enter your login information in the following format:

watchserver
4. In the User and Password boxes, enter the user name and password that will be used to access RenderPerfect. The default user name is **admin**. The default password is **123**.
5. Click **OK** to return to the Configure dialog box.
6. Click **OK** to return to Vista Configuration.

To edit a RenderPerfect login

1. Access the Configure dialog box.
2. Highlight, and the double-click the RenderPerfect login you want to modify. The Edit Login dialog box appears.

3. In the **What** box, enter your login information in the following format:
watchserver
4. In the **User** and **Password** boxes, enter the user name and password that will be used to access RenderPerfect.
5. In the **User** and **Password** boxes, enter the new user name and password that will be used to access RenderPerfect.
6. Click **OK** to return to the Configure dialog box.
7. Click **OK** to return to Vista Configuration.

To delete a RenderPerfect login

1. Access the Configure dialog box.
2. Highlight the login you want to delete.
3. Click the red X and the login is deleted.
4. Click **OK** to return to the Configure dialog box.
5. Click **OK** to return to Vista Configuration.

Uninstall Rendition Services on a Server

Uninstalling Rendition Services, removes all executable, plug-in and repository files from the computer's hard drive.

Note You can delete any remaining Rendition Services files, such as log files manually.

To uninstall Rendition Services on a server

1. From the Windows taskbar, access the Control Panel.
2. Double-click **Add/Remove Programs**.
3. In the Add/Remove Programs window, select **FileNet Rendition Server**.
4. Click **Remove**.
5. When you're prompted to confirm the removal, click **Yes**. A progress panel displays the status of the removal.
6. After the uninstall process is complete, reboot your computer and log in as the service account user.

Glossary

You should be familiar with the following terms and concepts. If you encounter other terms in the Content Services documentation that you do not understand, refer to the *CS Explorer Online Help*.

General Terms and Concepts

additional storage server A server where you installed additional services for a library system.

additional services A Storage Manager and a Content Search Manager installed as a unit on a server for a library system after initial services are installed for the library system on some other server.

application server A computer that provides library system services to client workstations. A typical server would contain one or more library system service modules and might be a workstation installed with UNIX or Windows software. Network operating system software may also be required in some network environments. See also network server.

archive category Describes a way of handling version files that are moved offline during automatic or manual archiving. Each protected item is given an archive category (for example, Archive To Tape) when it is added to the system. Each archive category is assigned to a specific storage directory called an archive repository.

archive repository A storage area (that is, a directory) for specific system-controlled version files that are moved offline. The archive repository is where the files are held while awaiting archive storage or destruction outside of the library system. Archive repositories can be described as reclaimable or nonreclaimable. Typically, files that are moved offline through a reclaimable archive repository are more easily reclaimed for use at a later time. Each archive repository may be assigned numerous archive categories.

Content Search Manager A software service that contains the program files required to make content searching for version files available with the library system.

content search repository A named collection of files identifying versions that have been indexed so that users can find those versions based on the actual contents of the version files. You must add at least one of these content search repositories for each Content Search Manager.

Database Custodial Server The storage server on which the daemons run to (1) update the database usage information (e.g., space calculations) and (2) delete outdated Content Services user sessions associated with the library system. By default, the Content Services Setup program sets this property value to the node name of the initial storage server. You may change this value by selecting the node name of another storage server associated with the library system. Such a change takes effect the next time the daemons run.

database server A server containing a database engine. If a library system uses that database engine, the database server is also called a *property server*.

initial services A Storage Manager, Content Search Manager, and Replication Services installed as a unit on a server for a library system.

initial storage server The server where you installed the initial services for a library system.

library system A set of initial services and zero or more sets of additional services

network server A computer that provides services for the network operating system under which the library system runs. Typically, the library system *application servers* are also network servers. See also *application server*.

Property Manager The components in the library system that store and maintain in a database a library system's objects and properties. These objects and properties identify and describe stored documents and unstored items. They also describe users, groups, searches, and system components. Each library system has its own Property Manager. The Property Manager is roughly equivalent to the database engine, but has a slightly larger conceptual purpose. For example, each of the library systems that happens to share a database engine is said to have its own unique Property Manager.

property repository A Microsoft SQL Server database or Oracle tablespace that contains the objects and properties for a single library system. The objects and properties identify and describe stored documents and unstored items. They also describe users, groups, searches, and system components.

The documentation generally refers to a SQL Server database or Oracle tablespace as a *property repository* only after you have installed the associated library system and applied its initial changes to the database or tablespace. Prior to library system installation, the documentation uses *database* or *tablespace* rather than *property repository* (for example, "To create an Oracle tablespace:").

property server The server where the library system database engine is installed. Typically, this server is also the initial storage server, but under certain configurations, may be a standalone database server instead.

server A computer that provides services for the network operating system or library system services. See also *network server* and *application server*.

service A unit of software that runs on a particular machine. Each library system service is typically referred to as a Manager, as in Property Manager, Storage Manager, and Content Search Manager. For example, every library system application server must be loaded with a Storage Manager service, which processes most of the file transfers and searches.

standalone database server A server where a database engine is installed and where no initial or additional services for any library system are installed.

Storage Manager The library system service that processes operational requests (for example, to check in or check out) and stores version files controlled by the library system.

storage category To a user, such a category typically describes a type of file (for example, Word Docs, Payroll Records, Memos). To an administrator, storage categories provide a way of controlling load balancing in the various storage repositories. Storage category names are determined by an administrator to reflect the working environment at a particular site. One or more storage categories may be assigned to a storage repository and administrators can change these assignments as load requirements change.

storage repository A storage area (that is, a directory) for specific version files controlled by the library system. Each file is stored according to a storage category, which is assigned when the file is added to the library system as a protected item. A storage repository may be assigned numerous storage categories for a particular installation of a library system.

storage server Any server where initial or additional services for a library system are installed.

Verity K2 Master Admin The Verity component that governs all Content Search Managers for a library system.

Verity K2 Master Admin Server The server where the CS installer places the Verity K2 Master Admin that governs the Content Search Managers for a given library system.

Workflow Daemon The software responsible for managing access and distribution of shared workflow components on a library system. (Applicable only if you have installed Panagon eProcess). See *Installing a Library System*.

Workflow Daemon user The *Workflow Daemon* uses the *Workflow Daemon* user to log in and out of the participating library systems to perform workflow operations. The *Workflow Daemon* user is assigned an encrypted password and has Administrator access rights to the specified library system. Do **not** modify this password. (Applicable only if you have installed Panagon eProcess). See *Installing a Library System*.

Replication-specific Terms and Concepts

administrative domain See *replication administrative domain*.

document replica A replica of an original (source) document.

Domain Master The library system where an *administrative domain* is created and maintained. For ease of maintenance, use the Domain Master to create and manage users, groups, Document Classes, and CVLs for the administrative domain.

folder replica A replica of an original folder.

home library system See *library system*. The originating library system for an object.

master replication library system See *original library system*.

original object The original object that is replicated to other libraries. For example, an original document or original folder. Modifications made to a original object are replicated to all replica objects automatically by the Replication Daemon. Original objects were previously called source objects.

original library system The library system where the original object is added. Original documents and original folders are created or added on the original library system. Original library systems were previously called source libraries.

ReplDaemon user The *Replication Daemon* uses the *ReplDaemon* user to log in and out of the participating library systems to perform replication operations. The *ReplDaemon* user is assigned an encrypted password and has Administrator access rights to the specified library system. Do **not** modify this password.

replica The replicated (copied) object that retains its link to its original (source) object, reflecting any changes that may occur in the original object. Also referred to as a replicated object or target object. For example, a folder replica or document replica. Modifications made to a replica are not always replicated back to the original object or any other replicated objects by the Replication Daemon.

replication administrative client The client used to administer the replication services for the clients.

replication administrative domain A collection of library systems that share security features. Within an administrative domain, any changes to User, Group, or Group Member objects are automatically replicated to the library systems. Stored searches created with the System Administration Tool, Document Classes, and CVLs are

also replicated across the domain. The initial user and initial user group (Admin and Administrators group) and their associated stored searches are not replicated.

Replication Daemon The replication software responsible for maintaining consistency between original objects and target objects.

replication folder Refers to a root-level folder that the administrator uses to designate where documents and subfolders are replicated. Once a folder is replicated, documents and subfolders can be added from any participating library. Participating libraries are listed in the folder's Comment field.

Replication Services The library system software and services that allow documents or folders to be replicated from one library system to another.

source document See *original object*.

source folder See *original object*.

source library system See *original library system*.

target library system The library system where replicas are located. The replica is created and owned by the *Replication Daemon* on the target library.

target replication folder A folder replica. See *Replication folder*.

A

- About the documentation 12
- Access Control
 - groups 33
 - in CS Explorer 34
 - multilevel security 20
- Access Domain
 - property 189
 - setting 194
- Access List
 - defaults 35
 - definition of 34
 - Item 35
 - replication folders 177
- Access Rights
 - Admin 33
 - affects on folder replicas 629
 - Author 33
 - description of 33
 - introduction 20
 - None 33
 - Owner 33
 - Viewer 33
- ACL
 - Replication Configuration 176
- Acrobat Plug-in options
 - defaults 696
 - setting 682
- Add Groups? property 264
- Add Items? property 264
- Add Users? property 264
- Added By User
 - in Group object 268
 - in User object 263
- Adding
 - a replication folder 623
 - archive categories 251
 - archive repository 247
 - content search repositories 255
 - custom properties 258, 259
 - group object 268
 - library systems to replication administrative domain 621
 - replication administrative domain 621
 - server-based configuration settings for Replication Service 624
 - storage categories 241
- Adding additional library systems 205
- Admin (access rights) 33
- Administrative Client
 - adding library systems to replication administrative domains 621
- Administrative Domain
 - configuring 621
 - considerations 598
 - establishing 621
 - overview 598
 - setting up 621
- Administrative domain
 - about 606
 - domain master 606
 - domain master, definition 720
 - restrictions 607
 - shared searches 606
- Administrators Group, customizing name of 27
- Adobe software
 - installing for Rendition Services 663
- AGGREGATE, Client Library Files setup command-line flag 209
- Annotations
 - replicating 612
- ANSI Character Set
 - ESPSansi font 701
 - Rendition Services 700
- ARCHDATE utility 371
- Architecture, object-oriented 24
- Archive Categories
 - adding 251
 - default archive category 250
 - determine archive/recalim requirements before adding 250
 - document replicas 177
 - links to archive repositories 250
 - modifying 251
 - overview of 28, 250
 - Reclaimable? property 250
- Archive Repositories
 - adding 247
 - backups 337
 - modifying 248
 - moving 248
 - overview of 28, 242
 - reclaiming 242
- ASCII Character Set
 - Courier font 702
 - ESPSascii font 700

- Rendition Services 700
 - ASCII files, limited ability to index 30
 - ASCII Plug-in
 - ESPSascii font 685
 - how options interact 685
 - options for Rendition Services 683
 - ASCII Transformation Options 683
 - Assigning
 - library systems to a replication folder 623
 - Audit Logging
 - analyzing log data 329
 - AUDit Event ID 325
 - bcp utility 329
 - clearing the log file 327
 - copying log data 328
 - Date Time 326
 - disk space requirements 324
 - Elapsed Time 326
 - Event Parameters 325
 - event parameters 326
 - events monitored 324
 - features 323
 - file format 325
 - Filename 326
 - Flags 326
 - implementing 324
 - increment set number 328
 - log file management 327
 - log file size 327
 - managing 324
 - managing the log file 327
 - Object ID 325
 - Object Key 325
 - overview 323
 - Param 1 325
 - parameters 1-4 326
 - Process ID 326
 - reports 329
 - sample Audit Log 326
 - Search Request Count 325
 - Search Request Direction 325
 - Search Request Start Key 325
 - Set Number 326
 - set number 328
 - set number, incrementing 328
 - setting up 324
 - Status 326
 - stpIncrementSet 328
 - stpTruncAuditByDate 328
 - stpTruncAuditLog 328
 - truncate by date 328
 - truncate by set number 328
 - turning on and off 325
 - User 326
 - viewing the audit log 325
 - Workstation 326
 - Author (access rights) 33
 - auto_archive_delay parameter
 - in initialization files 168
- B**
- Backups
 - archive repositories 337
 - content search repositories 338
 - housekeeping functions 337
 - planning 336
 - property repositories 337
 - Replication Services 603
 - storage repositories 337
 - Boolean searching, supported by content search 29
- C**
- CAT_INDX utility 375
 - CAT_PROC utility 377
 - CAT_STAT utility 379
 - CCAT utility 380
 - Character Sets
 - Rendition Services 700
 - Checking in and out document replicas 635
 - Clear definition of the DoD 331
 - client connectivity 112
 - Client Library Files
 - installing 208
 - installing or upgrading 207
 - installing unattended 208
 - setup 207
 - uninstalling 219
 - uninstalling unattended 220
 - Client Library Files Setup
 - silent mode 210
 - Clustering 38
 - Compound Documents
 - replicating 612
 - Configuration tool 164
 - Configuring
 - a replication folder 623
 - administrative domain 621

- Replication Services 618, 624
 - Configuring a library system, recommendations for 164
 - Connection pooling 108
 - Content Search
 - disable 254
 - enable 254
 - enabling 195
 - file filters supported for content searches 315
 - file formats supported 253
 - indexing daemon 253
 - installing 253
 - library system file filters 253
 - proximity searching 29
 - setting Default Index Versions? property 190
 - setting Default Indexed Version Limit property 193
 - stopwords 253
 - wildcard searching 29
 - Content Search Manager
 - directory created by library system 254
 - supported features for 29
 - Content Search Repositories
 - adding 255
 - backups 338
 - Max Repository Size property 254
 - modifying 256
 - moving 256
 - overview 253
 - Content Services user
 - creating 115
 - Content Types
 - default plug-in options 696
 - non-standard file extensions 682
 - plug-in options 682
 - Control panel services tool 201
 - Controlled Vocabulary Lists 366
 - replicating 606, 612
 - Controlling access using groups 267
 - COPY_CP utility 382
 - Corel Draw Plug-in
 - Rendition Services 687
 - Corel Presentations Plug-in
 - Rendition Services 687
 - Courier Character Set
 - Rendition Services 700
 - Courier font
 - Rendition Services 702
 - CS Configuration tool
 - Installation Tab 165
 - LDAP User Import Tab 180
 - Replication Tab 169
 - Storage Manager Tab 167
 - Windows User Import Tab 178
 - XML User Import Tab 182
 - CS Service Manager 198
 - CS system features
 - data caching 21
 - load balancing and scalability of distributed services 20
 - logical storage pointers 19
 - multi-level security 20
 - object-oriented approach 19
 - version control 20
 - CS_IMPORT utility 392
 - CSADMIN utility 384
 - CSFILTER utility 332
 - CSINDEX utility 332
 - CSMERGE utility 332, 338
 - Custom enterprise applications 20
 - Custom Properties
 - adding 258, 259
 - characteristics 258
 - integer value for 258
 - modifying 258, 261
 - multivalue 258
 - overview 36
 - overview of 258
 - setting up for replication 613
 - single-value 258
 - text value for 258
 - use with Document Classes and CVLs 366
 - CVLs
 - overview 366
 - replicating 606, 612
- ## D
- Daemon, replication 599
 - Data caching, description of 21
 - Database optimization 108
 - DBErrorLog parameter
 - in initialization files 168
 - DBHeartBeat parameter
 - in initialization files 167
 - DBThreshold parameter
 - in initialization files 167
 - Default Content Types
 - Rendition Services 696
 - Default values
 - Default Index Versions? 190

- Default Item Access List 190
 - Default Online Limit property 192
 - overview of how to set 188
 - selected from different properties 190
 - Definition
 - administrators group 267
 - objects 24
 - searches 29
 - Delete Upon Archive? property 242
 - Deleting
 - group object 268
 - user 264, 375
 - User object 264
 - Deleting items, minimum access level 195
 - Directory Server Type parameter
 - definition 166
 - Disabling
 - Replication Services 636
 - user logons 195
 - Display language 189, 196
 - Display Language property
 - in User object 266
 - overview 196
 - Document Classes
 - overview 366
 - replicating 606, 612
 - Document replicas
 - access control lists 634
 - checking in and out 635
 - creating 634
 - custom properties 613
 - definition 720
 - moving the original replicated document 635
 - overview 594
 - removing the original document 635
 - user and group access 629, 634, 646
 - Document Versions 25
 - Documentation updates 12
 - Documentation, about 12
 - Documentation, reader comments on 14
 - Documents
 - replicating 609
 - Domain Master
 - setting up an administrative domain 621
 - Domain master
 - administrative domain 606
 - definition 720
 - DSINFO utility 398, 636
 - dssetup
 - using with setup.inf 364
 - DSSTOP utility 400, 636
 - DSSTRT utility 402, 636
 - dsuninst.exe 218
 - DTCVLCHK utility 405
 - DTCVLMIG utility 412
- E**
- Enabled? property
 - in User object 265
 - overview 269
 - Enabling
 - Replication Services 618
 - Enterprise content management 19
 - Error Alert To
 - Replication Configuration 169
 - ErrorCode, Client Library Files Setup output file entry 212
 - ESPSansi font
 - Rendition Services 701
 - ESPSascii font
 - Rendition Services 700
 - Establishing an administrative domain 621
 - exceptions 628
- F**
- Feedback
 - on documentation 14
 - File filters
 - supported for content searches 30, 315
 - File filters supported for content searches 315
 - file format
 - Audit log 325
 - FILE2ID utility 419
 - Files, storage of 28
 - Firewalls, configuring for 117
 - Folder icon for replicas 623
 - Folder replicas
 - access rights 629, 646
 - creating 629, 646
 - creating from existing folders 629
 - definition 720
 - managing 629, 646
 - modified by daemon 629, 646
 - moving 631
 - removing replication folders 631
 - Folders
 - replicating 609

- FolderTitleFormat
 - Replication Configuration 175
- Font Characters
 - Rendition Services 700
- FTP site, for software patches 12
- FullAction
 - Replication Configuration 170
- FullAlertTo
 - Replication Configuration 170

G

- General Users Group, customizing name of 27
- Glossary
 - general terms 718
 - replication terms 720
- Graphic Files Plug-in
 - Rendition Services 687
- Group defaults, imported 184
- Group Object
 - adding 268
 - controlling search results 270
 - default values in 269
 - Enabled? property 269
 - modifying 269
 - modifying access list 269
 - modifying administrator's group 270
 - overview 26, 267
 - Owner access rights to 269
 - properties modified with Admin access rights 268
 - renaming 268
 - setting up administrator's group 270
- Groups
 - access rights determined by 33
 - active 34
 - adding 267
 - adding users to 267
 - authority to add 264
 - deleting 268
 - modifying properties 269

H

- HMSRVCHK utility 420
- HMSRVMIG utility 424, 626
- Home library system
 - see Original library system 720
- hosts file 216

I

- ID2FILE utility 428
- Image Services, installing with Content Services 80
- Importing Users and Groups 272
 - Usage Guidelines 285
 - Windows User Import Log File Example 296
 - Windows User Import Usage Example 294
- IN, Client Library Files setup command-line flag 209
- Indexing
 - limited ability to index ASCII files 30
 - replicas 172
 - stopwords 320
- IndexReplicas
 - Replication configuration 172
- INDXDATE utility 429
 - run from a batch file 431
 - run from command line 430
- INDEXFIX utility 432
- Initialization file
 - auto_archive_delay 168
 - DBErrorLog 168
 - DBHeartBeat 167
 - DBThreshold 167
 - Log File 167
 - Max File Size 167
 - ReclaimCheck 168
 - Secure Deletes 168
- Installation
 - on Windows 2000 servers 123
 - prerequisites 80
 - roadmap 68
 - troubleshooting problems 215
- InstallFileVersion, Client Library Files Setup input file option 210
- Installing
 - a library system unattended 136
 - accessing multiple database types from one client 206
 - adding additional library systems 205
 - additional library system components 130
 - Client Library Files 207
 - Content Services with Image Services 80
 - IDMDS 119
 - Non-English EnvironmentsS 81
 - Oracle8i Client files 109, 111
 - problem solving 215
 - Rendition Services software 664
 - with setup.inf 364
- Installing an iPlanet LDAP Server 276

- Insufficient disk space 215
- Insufficient memory 215
- Internal Custom Properties 258
- Item access lists 35
- Item properties
 - Replication Services 615
- Items
 - protected, definition of 24
 - unprotected, definition of 25
 - versions associated with 25
- ItemTitleFormat
 - Replication Configuration 174
- L**
- Language
 - assigning for display 189, 196
 - changing system language 196
- LDAP_EXP utility 435
- Library
 - Replication Configuration 171
- Library system
 - adding to replication administrative domain 621
 - administrative domain 621
 - configuration settings for 164
 - domain master, definition 720
 - error logs for 215
 - installing Client Library Files for 207
 - installing with setup.inf 364
 - maintaining 164
 - multiple with one database engine 204
 - objects supported for replication 597
 - problem solving during Setup of 215
 - recommendations for configuring and maintaining 164
 - server unresponsive to network communication 216
 - unattended uninstall 218
 - uninstalling 217
 - upgrading 152
- library system 22
- Library system events
 - events in audit log 324
- Library system naming requirements 82
- Library system planning guide 39
- Library system services
 - controlling 198
 - controlling with command-line utilities 202
 - controlling with Service Manager 198
 - controlling with Windows control panel 201
- Links
 - archive categories to archive repositories 242
 - storage categories to storage repositories 234
- Load balancing, description of 20
- LOCATION, Client Library Files setup command-line flag 209
- Log File parameter
 - in initialization files 167
- log_buffer parameter 108
- Logical storage pointers, description of 19
- Logon
 - customizing 195
 - disabling user logons 195
 - Replication Configuration 170
 - restricting to Administrators only 195
- Logons Enabled? property 195
- Logs
 - for Audit Logging 327
 - for installation errors 215
- Lost+found directory 332
- M**
- MacroMedia Freehand Plug-in
 - Rendition Services 688
- Managing
 - Document Classes and CVLS in Replication Services 600
 - folder replicas 629, 646
 - Replication Services 629
- Max File Size parameter
 - in initialization files 167
- MaxSize
 - Replication Configuration 169
- Memory
 - insufficient 215
 - responding to insufficient RAM error message 215
- Minimum item delete access 195
- Modifying
 - archive categories 251
 - archive repositories 248
 - content search repositories 256
 - custom properties 258, 261
 - group object access list 269
 - Group properties 269
 - user 263
- Modifying Replicas
 - inside an administrative domain 632
 - outside an administrative domain 632
 - overview 632
- Moving

- archive repositories 248
 - content search repositories 256
 - folder replicas 631
 - original replicated documents 635
 - MS PowerPoint Plug-in
 - Rendition Services 689
 - MS Project Plug-in
 - Rendition Services 690
 - MS Word Plug-in
 - Rendition Services 691
 - Multi-level security, description of 20
 - Multiple library systems
 - accessing with one client 206
 - common configuration features 310
 - overview of 310
 - single database engine 204
 - MustReboot, Client Library Files Setup output file entry 212
- N**
- Names of replicas 174, 175
 - Names, custom 258
 - Naming requirements
 - library system 82
 - server 81
 - tablespace 82
 - net service name configuration 112
 - Network adapter cards 216
 - Network logon Name property 265
 - None (access rights) 33
- O**
- Object-oriented
 - approach, description of 19
 - architecture 24
 - Objects
 - definition of 24
 - Group 26
 - User 26
 - Version 25
 - ODBC connection for Oracle9i 112
 - Online Version Limit property for items 192
 - Optimizing the database 108
 - Oracle configuration
 - optimization 108
 - Oracle parameters
 - log_buffer 108
 - processes 108
 - shared_pool_size 108
 - Oracle tablespace
 - creating 106
 - Oracle8i Client files
 - installing 109, 111
 - Oracle9i Client Connectivity 112
 - Oracle9i net service name configuration 112
 - Oracle9i ODBC Connection 112
 - Oracle9i tablespace
 - creating 102, 104
 - Original library system
 - definition 720
 - original library system
 - Replication Services 604
 - Original object, definition 720
 - OUT, Client Library Files setup command-line flag 209
 - Overview
 - archive categories 28, 250
 - archive repositories 28, 242
 - custom properties 258
 - CVLs 366
 - Document Classes 366
 - group objects 267
 - Library System 366
 - multiple library systems 310
 - replicas 594
 - Replication Services 594
 - storage categories 28, 240
 - storage repositories 28, 233
 - user object 263
 - Owner (access rights) 33
- P**
- Passwords
 - for imported users 284
 - Passwords for Imported Users 284
 - Patches, FTP site 12
 - Pausing
 - Replication Services 636
 - Permanent Versions in Limit? property 192
 - Planning
 - Replication Services 601
 - Planning guide 39
 - Plug-in
 - ASCII options, Rendition Services 683
 - Corel Draw options, Rendition Services 687
 - Corel Presentations options, Rendition Services 687

- default content types 696
 - default plug-in options 696
 - Graphic Files options, Rendition Services 687
 - how ASCII options interact 685
 - Macromedia Freehand options, Rendition Services 688
 - MS PowerPoint options, Rendition Services 689
 - MS Project options, Rendition Services 690
 - MS Word options, Rendition Services 691
 - non-standard file extensions 682
 - PowerPoint options, Rendition Services 689
 - Project options, Rendition Services 690
 - setting content types 682
 - setting plug-in options 682
 - text file options, Rendition Services 683
 - Word options, Rendition Services 691
 - WordPerfect options, Rendition Services 688
 - Plug-ins
 - choosing a plug-in 681
 - Port number
 - Content Search Manager 117
 - Power interruption
 - effect on installation 216
 - PowerPoint Plug-in
 - Rendition Services 689
 - Prerequisites 80
 - Rendition Services 659
 - Replication Services 618
 - Preview transaction file, User Import tool 307
 - Primary Storage Manager property
 - in System Object 191
 - in User object 266
 - Privileges of access 33
 - Problem solving 215
 - for library system installation 215
 - processes parameter 108
 - Program defaults for access list 190
 - Project Plug-in
 - Rendition Services 690
 - Properties, custom 24
 - Property
 - Add Groups? 264
 - Add Items? 264
 - Add Users? 264
 - Delete Upon Archive? 242
 - Display Language 266
 - Enabled? 265, 269
 - Max Repository Size 254
 - Network logon Name 265
 - Primary Storage Manager 266, 269
 - Reclaimable? 250
 - Property repository
 - restore operation settings 339
 - Property repository tablespace
 - creating 102, 104, 106
 - Protected Items 24
 - Proximity searching, supported by content search 29
 - Published Documents
 - replicating 612
 - Purge definition of the DoD 331
- ## R
- Reclaim daemon, changing the interval for 168
 - Reclaiming versions
 - process for 244
 - process request 245
 - properties modified by library system 245
 - Release notes 17
 - addendum 16
 - Remapping special characters
 - Rendition Services 698
 - RemoteIntervalMinutes
 - Replication Configuration 171
 - Removing
 - client files 219
 - library system 217
 - original documents 635
 - replication folders 631
 - Removing unattended
 - client files 220
 - Renaming
 - group object 268
 - user 264
 - Renaming Servers 348
 - Rendition Services 658, 659
 - ANSI character set 700
 - ASCII character set 700
 - ASCII plug-in 683
 - character sets 700
 - choosing plug-ins 681
 - Corel Draw Plug-in 687
 - Corel Presentations Plug-in 687
 - Corel WordPerfect Plug-in 688
 - courier character set 700
 - Courier font 702
 - default content types 696
 - default plug-in options 696
 - ESPSansi font 701

- ESPSascii font 685, 700
- font characters 700
- Graphic Files Plug-in 687
- how ASCII options interact 685
- installing on server 664
- MacroMedia Freehand Plug-in 688
- Microsoft Word, preparing documents 677
- MS PowerPoint Plug-in 689
- MS Project Plug-in 690
- MS Word Plug-in 691
- non-standard content types 682
- other Adobe software 663
- PowerPoint Plug-in 689
- prerequisites 659
- Project Plug-in 690
- remapping special characters 698
- requirements for installing 659
- setting content types 682
- setting plug-in options 682
- setting up on server 664
- text file plug-in 683
- Word Plug-in 691
- WordPerfect Plug-in 688
- WordPerfect printer 678
- WordPerfect, preparing documents 678
- REPCHECK utility 439
 - used during Secure Document Delete 332
- ReplDaemon user
 - access rights assigned to 620
 - applying Default Item Access List to replicas 176, 177
 - automatic creation of 620
 - definition 720
- repli.ini file
 - See CS Configuration Tool Replication Tab 169
- Replica titles 174, 175
- ReplicaIndexLimit
 - Replication Configuration 173
- ReplicaOnlineLimit
 - Replication Configuration 172
- Replicas
 - annotations 612
 - compound document 612
 - creating documents 634
 - creating folders 629, 646
 - custom properties 613
 - CVL 612
 - definition 720
 - deleting users and groups 634
 - document 609
 - Document Class 612
 - existing folders 629
 - folder 609
 - folder access rights 629, 646
 - folder icon 623
 - folders 629, 646
 - folders modified at creation 629, 646
 - icon for folders 629, 646
 - indexing 172
 - item properties 615
 - managing document replicas 634
 - managing folders 629, 646
 - modifying inside an administrative domain 632
 - modifying outside an administrative domain 632
 - modifying users and groups 633
 - overview 594
 - Published Documents 612
 - removing replication folders 631
 - stored searches 614
 - synchronizing with original objects 594
 - titles 174
 - user and group access 629, 634, 646
 - users and groups 614, 633
 - version properties 615
- ReplicateArchiveCategory
 - Replication Configuration 177
- ReplicateFolderACLs
 - Replication Configuration 177
- ReplicateItemACLs
 - Replication Configuration 176
- ReplicateStorageCategory
 - Replication Configuration 177
- Replicating
 - annotations 612
 - compound documents 612
 - custom properties 613
 - CVLs 612
 - Document Class 612
 - documents 609
 - folders 609
 - item and version properties 615
 - Published Documents 612
- Replication
 - access rights for folder replicas 629, 646
 - adding a replication folder 623
 - configuring 624
 - configuring a replication folder 623
 - creating folder replicas 629, 646

- defined 594
- definitions 595
- designating a replication folder 623
- document status 594
- existing folders 629
- folder replicas modified by daemon 629, 646
- glossary 720
- icons for folder replicas 629, 646
- identifying a replica 174, 175
- indexing replicas 172
- moving folder replicas 631
- objects supported for replication 597
- overview 594
- prerequisites 618
- removing replication folders 631
- ReplDaemon user 599
- ReplDaemon user, creation of 620
- synchronization of replicas 171
- terminology 595
- user and group access list entries for document replicas 629, 646
- view-only documents 599
- Replication administrative client, definition of 720
- Replication administrative domain
 - definition of 721
- Replication Configuration
 - Error Alert To 169
 - FolderTitleFormat 175
 - FullAction 170
 - FullAlertTo 170
 - ItemTitleFormat 174
 - Library 171
 - Logons 170
 - MaxSize 169
 - RemoteIntervalMinutes 171
 - ReplicaIndexLimit 173
 - ReplicaOnlineLimit 172
 - ReplicateArchiveCategory 177
 - ReplicateFolderACLs 177
 - ReplicateItemACLs 176
 - ReplicateStorageCategory 177
 - Requests 170
 - Run Times 177
 - SPI 171
 - Startup 169
 - Targets configuration parameter 171
 - Truncate To 169
- Replication configuration
 - IndexReplicas 172
 - WakeupSeconds 169
- Replication daemon
 - definition 721
 - definition of 721
 - overview 599
- Replication Folder 609
 - overview 609
- Replication folder
 - creating 623
 - definition 721
 - removing 631
- Replication Home Server 599
- Replication Home Server feature 599
- Replication Services 593, 606
 - annotations 612
 - backup and restore operations 603
 - checking in and out document replicas 635
 - compound documents 612
 - configuring 618
 - considerations 601
 - Controlled Vocabulary Lists 606
 - controlling 636
 - creating document replicas 634
 - custom properties 613
 - CVL 606, 612
 - definition 721
 - disabling 636
 - distribution of documents 602
 - Document Class 612
 - Document Classes 606
 - enabling 618
 - impact on library systems 601
 - impact on resources 603
 - item properties 615
 - managing 629
 - managing Document Classes and CVLS 600
 - moving original replicated documents 635
 - original library system 604
 - planning 601
 - Published Documents 612
 - removing original documents 635
 - replicating documents 609
 - replicating users and groups 633
 - Replication Folder 609
 - replication folder 623
 - replication folders 609
 - setting up 618
 - starting a service 636
 - stopping a service 636

- stored searches 614
 - tnsnames.ora file 621
 - user and group access list entries for document
 - replicas 634
 - users and groups 614
 - version properties 615
- Replication Utilities
 - DTCVLCHK 405
 - DTCVLMIG 412
 - REPLKICK 443
 - USGRPCHK 448
 - USGRPMIG 455
- REPLKICK utility 443
- Requests
 - Replication Configuration 170
- Requirements
 - Rendition Services 659
- Restore operation
 - settings for restoring property repository 339
- Restoring
 - procedure 339
- Restricting logons to administrators only 195
- ReturnCode, Client Library Files Setup output file entry 212
- Roadmap to setting up 68
- Run Times
 - Replication Configuration 177
- S**
- Sample XML File 297
- Sanitize definition of the DoD 331
- Scalability of distributed services, description of 20
- Scrubbing a file, definition of 331
- Search Domain property
 - in Group objects 270
 - in User object 266
- Search, shared
 - administrative domain 606
- Searches
 - controlling access to results of 193
 - definition of 29
 - group access, affects on results of 194
 - replicating 597, 609
 - replicating stored searches 614
 - setting Viewer security on results of 194
- Searching for file contents, See Content Search Manager 29
- Secure Deletes parameter
 - definition 168
- Secure Document Delete
 - Clear definition of the DoD 331
 - levels 331
 - overview 331
 - scrubbing 331
- Secure Searches? property 189, 194
- Secure Searches? property, in Group object 270
- Secure Searches? property, in User object 266
- Server
 - Replication Services configuration parameters 624
- Server naming requirements 81
- server, renaming 348
- Servers
 - delaying automatic version archiving on 168
- Service Manager 198
- Services
 - controlling 198
 - controlling with command-line utilities 202
 - controlling with Service Manager 198
 - controlling with Windows control panel 201
- Setting Content Types
 - Rendition Services 682
- Setting Plug-in options
 - Rendition Services 682
- Setting up
 - administrative domain 621
 - multiple library systems 204
 - Rendition Services software 664
 - Replication Services 618
- setup.inf
 - creating 364
 - modifying 364
 - using with dssetup 364
- Shared searches
 - administrative domain 606
- shared_pool_size parameter 108
- Silent mode
 - Client Library Files Setup 210
- Software patches, FTP site 12
- Source document, definition 720
- Source library system
 - definition 720
- Special characters
 - Rendition Services 698
- SPI
 - Replication Configuration 171
- SPISTAT utility 445
- SQL*Plus
 - using to create a tablespace 106

- SRDELETE utility 446
- Starting
 - library system services 198
 - Replication Services 636
- Startup
 - Replication Configuration 169
- Status codes
 - severity levels 573
 - table for 573
- Stopping
 - library system services 198
 - Replication Services 636
- Stopwords 29, 320
- Storage categories
 - adding 241
 - assigning to content search repositories 254
 - document replicas 177
 - links to content search repositories 253
 - links to storage repositories 233
 - overview of 28, 240
 - properties 240, 243
- Storage Manager
 - fails to bind socket 216
- Storage repositories
 - backups 337
 - overview of 28, 233
 - restore procedure 343
- stpIncrementSet 328
- stpTruncAuditByDate 328
- stpTruncAuditLog 328
- Synchronizing library system groups and users with Windows domain and LDAP groups and users 272, 292
- Synchronizing replicas 171, 594
- System features, CS
 - data caching 21
 - load balancing and scalability of distributed services 20
 - logical storage pointers 19
 - multi-level security 20
 - object-oriented approach 19
 - version control 20
- System object, modifying 164
- System planning
 - describing your 42
 - designing 42
 - estimating hardware and software requirements 42
- Systems Management Server 220

T

- Tablespace
 - creating 102, 104, 106
 - naming requirements 82
- Target library system, definition 721
- Target replication folder, definition 720, 721
- TargetDirectory, Client Library Files Setup input file option 211
- Targets
 - Replication Configuration 171
- Terminology
 - replication 595
- Text File Plug-in 683
- Text File Transformation Options 683
- tnsnames.ora file
 - Replication Services 621
 - specifying library systems in 621
- Tracking versions 31
- Transformation Options
 - ASCII plug-in 683
 - text file plug-in 683
- Troubleshooting installation problems 215
- Troubleshooting, See Problem solving 215
- Truncate To
 - Replication Configuration 169

U

- Unattended mode
 - Client Library Files Setup 210
- Uninstalling 217
 - Client Library Files 219
 - library system 217
 - Replication Services 636
 - Windows 2000-based library system 217
- Uninstalling unattended
 - Client Library Files 220
- Unprotected items 25
- Upgrade
 - library system 152
 - preparing for 154, 156
- Upgrading
 - Client Library Files 207
- User Import tabs 272
- User Import tool
 - imported group defaults 184
 - passwords 284
 - preview transaction file 307
- User Object

- assigning and modifying values 264
- deleting 264, 375
- modifying 263
- modifying access list 265
- overview 26
- overview of 263
- Owner access rights to 265
- properties for controlling user search results 266
- properties modified with Admin access rights 263
- renaming 264
- setting logon values 265
- Users
 - authority to add 264
 - deleting 264, 375
 - making members of groups 263
 - modifying 263
 - modifying properties 264
 - renaming 264
- Users and Groups
 - access list entries for document replicas 629, 634, 646
 - replicating 614, 633
- USGRPCHK utility 448
- USGRPMIG utility 455
- Using LDAP User Import 277
- Using Windows User Import 273
- Using XML User Import 281
- Utilities
 - ARCHDATE 371
 - CAT_INDX 375
 - CAT_PROC 377
 - CAT_STAT 379
 - CCAT 380
 - COPY_CP 382
 - CS_IMPORT 392
 - CSADMIN 384
 - CSMERGE 338
 - DSINFO 398
 - DSSTOP 400
 - DSSTRT 402
 - DTCVLCHK 405
 - DTCVLMIG 412
 - FILE2ID 419
 - ID2FILE 428
 - INXDDATE 429
 - INDXFIX 432
 - LDAP_EXP 435
 - overview 369
 - REPCHECK 439

- REPLKICK 443
- SPISTAT 445
- SRDELETE 446
- USGRPCHK 448
- USGRPMIG 455
- VUPDATE 461

V

- Variants file
 - supported by content search 30
- Version control, description of 20
- Version line, definition of 32
- Version Online Limits 191, 193
- Version properties
 - Replication Services 615
- Versions
 - document 25
 - manual archiving 243
 - nonreclaimable 242
 - object 25
 - tracking 31
- Viewer (access rights) 33
- view-only documents
 - replication 599
- VUPDATE utility 461

W

- Wakeup
 - replication configuration 169
- WANs, configuring for 117
- Warnings, Client Library Files Setup output file entry 212
- Wildcard searching supported for content search 29
- Windows 2000
 - installing a library system 123
 - uninstalling a library system 217
- Word Plug-in
 - Rendition Services 691
- WordPerfect Plug-in
 - Rendition Services 688
- WordPerfect Printer
 - Rendition Services 678

X

- XML User Import using LDIF (LDAP Data Interchange Format) Format Files 289

